

(19) World Intellectual Property
Organization
International Bureau



(43) International Publication Date
29 September 2005 (29.09.2005)

PCT

(10) International Publication Number
WO 2005/091593 A1

(51) International Patent Classification⁷: **H04L 29/06**,
12/28, 12/56

(21) International Application Number:
PCT/US2005/008832

(22) International Filing Date: 17 March 2005 (17.03.2005)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
60/554,309 17 March 2004 (17.03.2004) US
60/556,345 24 March 2004 (24.03.2004) US

(71) Applicant (for all designated States except US): QUALCOMM Incorporated [US/US]; 5775 Morehouse Drive, San Diego, California 92121 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): ANDERSON, Jon James [US/US]; 7436 Augusta Drive, Boulder, Colorado

80301 (US). STEELE, Brian [US/US]; 1074 Iliad Way, Lafayette, Colorado 80026 (US). WILEY, George A. [US/US]; 5740 Brittany Forrest Lane, San Diego, California 92130 (US). SHEKHAR, Shashank [IN/US]; 1597 Holman Drive, Erie, Indiana 80516 (US).


(74) Agents: WADSWORTH, Philip R. et al.; 5775 Morehouse Drive, San Diego, California 92121 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GI, GM, GR, GU, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MY, NA, NI, NO, NZ, OM, PG, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, YU, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM),

[Continued on next page]

(54) Title: HIGH DATA RATE INTERFACE APPARATUS AND METHOD



(57) Abstract: A data interface for transferring digital data between a host and a client over a communication path using packet structures linked together to form a communication protocol for communicating a pre-selected set of digital control and presentation data. The signal protocol is used by link controllers configured to generate, transmit, and receive packets forming the communications protocol, and to form digital data into one or more types of data packets, with at least one residing in the host device and being coupled to the client through the communications path. The interface provides a cost-effective, low power, bi-directional, high-speed data transfer mechanism over a short-range "serial" type data link, which lends itself to implementation with miniature connectors and thin flexible cables which are especially useful in connecting display elements such as wearable micro-displays to portable computers and wireless communication devices.



European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LI, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IS, IT, LI, LU, MC, NL, PL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG)

Decisions under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii)) for the following designations AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NA, NI, NO, NZ, OM, PG, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, SM, SY, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW, ARIPO patent (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SI, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT,

- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii)) for all designations

Published:

- with international search report
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments

For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.

HIGH DATA RATE INTERFACE APPARATUS AND METHOD

Claim of Priority under 35 U.S.C. §119

[0001] The present Application for Patent claims priority to Provisional Application No. 60/554,309 entitled "Switchable Threshold Differential Interface" filed March 17, 2004, and No. 60/556,345 entitled "Switchable Threshold Differential Interface" filed March 24, 2004, both applications of which are assigned to the assignee hereof and hereby expressly incorporated by reference herein.

BACKGROUND

Field

[0002] Embodiments of the invention in this disclosure relate to a digital signal protocol, process, and apparatus including integrated circuits and components for communicating or transferring signals between a host device and a client device at high data rates. More specifically, the disclosure relates to a technique for transferring multimedia and other types of digital signals from a host or controller device to a client device for presentation or display to an end user using a low power high data rate transfer mechanism having internal and external device applications.

Background

[0003] Computers, electronic game related products, and various video technologies (for example Digital Versatile Disks (DVD's) and High Definition VCRs) have advanced significantly over the last few years to provide for presentation of increasingly higher resolution still, video, video-on-demand, and graphics images, even when including some types of text, to end users of such equipment. These advances in turn mandated the use of higher resolution electronic viewing devices such as high definition video monitors, High Definition Television (HDTV) monitors, or specialized image projection elements. Combining such visual images with high-definition or -quality audio data, such as when using Compact Disc (CD) type sound reproduction, DVDs, surround-sound, and other devices also having associated audio signal outputs, is used to create a more realistic, content rich, or true multimedia experience for an end user. In addition, highly mobile, high quality sound systems and music transport mechanisms,

such as MP3 players, have been developed for audio only presentations to end users. This has resulted in increased expectations for typical users of commercial electronics devices, from computers to televisions and even telephones, now being accustomed to and expecting high or premium quality output.

[0004] In a typical video presentation scenario, involving an electronics product, video data is typically transferred using current techniques at a rate that could be best termed as slow or medium, being on the order of one to tens of kilobits per second. This data is then either buffered or stored in transient or longer-term memory devices, for delayed (later) play out on a desired viewing device. For example, images may be transferred "across" or using the Internet using a program resident on a computer having a modem or other type of Internet connection device, to receive or transmit data useful in digitally representing an image. A similar transfer can take place using wireless devices such as portable computers equipped with wireless modems, or wireless Personal Data Assistants (PDAs), or wireless telephones.

[0005] Once received, the data is stored locally in memory elements, circuits, or devices, such as Random Access Memory (RAM) or flash memory, including internal or external storage devices such as small size hard drives, for playback. Depending on the amount of data and the image resolution, the playback might begin relatively quickly, or be presented with longer-term delay. That is, in some instances, image presentation allows for a certain degree of real time playback for very small or low resolution images not requiring much data, or using some type of buffering, so that after a small delay, some material is presented while more material is being transferred. Provided there are no interruptions in the transfer link, or interference from other systems or users relative to the transfer channel being used, once the presentation begins the transfer is reasonably transparent to the end user of the viewing device. Naturally, where multiple users share a single communication path, such as a wired Internet connection, transfers can be interrupted or slower than desired.

[0006] The data used to create either still images or motion video are often compressed using one of several well known techniques such as those specified by the Joint Photographic Experts Group (JPEG), the Motion Picture Experts Group (MPEG), and other well known standards organizations or companies in the media, computer, and communications industries to speed the transfer of data over a communication link.

This allows transferring images or data faster by using a smaller number of bits to transfer a given amount of information.

[0007] Once the data is transferred to a "local" device such as a computer having a storage mechanism such as memory, or magnetic or optical storage elements, or to other recipient devices, the resulting information is un-compressed (or played using special decoding players), and decoded if needed, and prepared for appropriate presentation based on the corresponding available presentation resolution and control elements. For example, a typical computer video resolution in terms of a screen resolution of X by Y pixels typically ranges from as low as 480x640 pixels, through 600x800 to 1024x1024, although a variety of other resolutions are generally possible, either as desired or needed.

[0008] Image presentation is also affected by the image content and the ability of given video controllers to manipulate the image in terms of certain predefined color levels or color depth (bits per pixel used to generate colors) and intensities, and any additional overhead bits being employed. For example, a typical computer presentation would anticipate anywhere from around 8 to 32, or more, bits per pixel to represent various colors (shades and hues), although other values are encountered.

[0009] From the above values, one can see that a given screen image is going to require the transfer of anywhere from 2.45 Megabits (Mb) to around 33.55 Mb of data over the range from the lowest to highest typical resolutions and depth, respectively. When viewing video or motion type images at a rate of 30 frames per second, the amount of data required is around 73.7 to 1,006 Megabits of data per second (Mbps), or around 9.21 to 125.75 Megabytes per second (MBps). In addition, one may desire to present audio data in conjunction with images, such as for a multimedia presentation, or as a separate high resolution audio presentation, such as CD quality music. Additional signals dealing with interactive commands, controls, or signals may also be employed. Each of these options adding even more data to be transferred. Furthermore, newer transmission techniques involving HDTV and movie recordings may add even more data and control information. In any case, when one desires to transfer high quality or high resolution image data and high quality audio information or data signals to an end user to create a content rich experience, a high data transfer rate link is required between presentation elements and the source or host device that is configured to provide such types of data.

[0010] Data rates of around 115 Kilobytes (KBps) or 920 Kilobits per second (Kbps) can be routinely handled by some modern serial interfaces. Other interfaces such as USB serial interfaces can accommodate data transfers at rates as high as 12 MBps, and specialized high speed transfers such as those configured using the Institute of Electrical and Electronics Engineers (IEEE) 1394 standard, can occur at rates on the order of 100 to 400 MBps. Unfortunately, these rates fall short of the desired high data rates discussed above which are contemplated for use with future wireless data devices and other services for providing high resolution, content rich, output signals for driving portable video displays or audio devices. This includes computers for business and other presentations, gaming devices, and so forth. In addition, these interfaces require the use of a significant amount of host or system and client software to operate. Their software protocol stacks also create an undesirably large amount of overhead, especially where mobile wireless devices or telephone applications are contemplated. Such devices have severe memory and power consumption limitations, as well as already taxed computational capacity. Furthermore, some of these interfaces utilize bulky cables which are too heavy and unsatisfactory for highly aesthetic oriented mobile applications, complex connectors which add cost, or simply consume too much power.

[0011] There are other known interfaces such as the Analog Video Graphics Adapter (AVGA), Digital Video Interactive (DVI) or Gigabit Video Interface (GVIF) interfaces. The first two of these are parallel type interfaces which process data at higher transfer rates, but also employ heavy cables and consume large amounts of power, on the order of several watts. Neither of these characteristics are amenable to use with portable consumer electronic devices. Even the third interface consumes too much power and uses expensive or bulky connectors.

[0012] For some of the above interfaces, and other very high rate data systems/protocols or transfer mechanisms associated with data transfers for fixed installation computer equipment, there is another major drawback. To accommodate the desired data transfer rates also requires substantial amounts of power and/or operation at high current levels. This greatly reduces the usefulness of such techniques for highly mobile consumer oriented products.

[0013] Generally, to accommodate such data transfer rates using alternatives such as say optical fiber type connections and transfer elements, also requires a number of additional converters and elements that introduce much more complexity and cost, than

desired for a truly commercial consumer oriented product. Aside from the generally expensive nature of optical systems as yet, their power requirements and complexity prevents general use for lightweight, low power, portable applications.

[0014] What has been lacking in the industry for portable, wireless, or mobile applications is a technique to provide a high quality presentation experience, whether it be audio, video, or multimedia based, for highly mobile end users. That is, when using portable computers, wireless phones, PDAs, or other highly mobile communication devices or equipment, the current video and audio presentation systems or devices being used simply cannot deliver output at the desired high quality level. Often, the perceived quality that is lacking is the result of unobtainable high data rates needed to transfer the high quality presentation data. This can include both transfer to more efficient, advanced or feature laden external devices for presentation to an end user, or transfer between hosts and clients internal to portable devices such as computers, gaming machines, and wireless devices such as mobile telephones.

[0015] In this latter case, there have been great strides made in adding higher and higher resolution internal video screens, and other specialty input and/or output devices and connections to wireless devices like so called third generation telephones, and to so called laptop computers. However, internal data busses and connections which may include bridging across rotating or sliding hinge or hinge-like structures which mount or connect video screens or other elements to the main housing where the host and/or various other control elements and output components reside. These are generally high-bandwidth or high throughput interfaces. It is very difficult to construct high throughput data transfers interfaces using prior techniques which can require up to 90 conductors, or more, to achieve the desired throughput, on say a wireless telephone, as one example. Current solutions typically involve employing parallel type interfaces with relatively high signal levels which can cause the interconnection to be more costly, less reliable and potentially generate radiated emissions which could interfere with device functions. This presents many manufacturing, cost, and reliability challenging issues to overcome.

[0016] Such issues and requirements are also being seen on fixed location installations where communication or computing type devices, as one example, are added to appliances and other consumer devices to provide advanced data capabilities, Internet and data transfer connections, or built in entertainment. Another example would be

airplanes and busses where individual video and audio presentation screen are mounted in seat-backs. However, in these situations it is often more convenient, efficient, and easily serviceable to have the main storage, processing, or communication control elements located a distance from visible screens or audio outputs with an interconnecting link or channel for the presentation of information. This link will need to handle a significant amount of data to achieve the desired throughput, as discussed above.

[0017] Therefore, a new transfer mechanism is needed to increase data throughput between host devices providing the data and client display devices or elements presenting an output to end users.

[0018] Applicants have proposed such new transfer mechanisms in U.S. Patent Application Serial No. 10/020,520, filed December 14, 2001, now U.S. Patent No. 6,760,772, issued July 6, 2004 to Zou et al., and U.S. Patent Application Serial No. 10/236,657, filed September 6, 2002, both entitled "Generating And Implementing A Communication Protocol And Interface For High Data Rate Signal Transfer," both of which are assigned to the assignee of the present invention and incorporated herein by reference. Also, U.S. Application Serial No. 10/860,116, filed on June 2, 2004, entitled "Generating and Implementing a Signal Protocol and Interface for Higher Data Rates." The techniques discussed in those applications can greatly improve the transfer rate for large quantities of data in high speed data signals. However, the demands for ever increasing data rates, especially as related to video presentations, continue to grow. Even with other ongoing developments in data signal technology, there is still a need to strive for even faster transfer rates, improved communication link efficiencies, and more powerful communication links. Therefore, there is a continuing need to develop a new or improved transfer mechanism which is needed to increase data throughput between host and client devices.

SUMMARY

[0019] The above drawback, and others, existent in the art are addressed by embodiments of the invention in which a new protocol and data transfer means, method and mechanism have been developed for transferring data between a host device and a recipient client device at high data rates.

[0020] Embodiments for the invention are directed to a Mobile Data Digital Interface (MDDI) for transferring digital data at a high rate between a host device and a client device over a communication path which employs a plurality or series of packet structures to form a communication protocol for communicating a pre-selected set of digital control and presentation data between the host and client devices. The signal communications protocol or link layer is used by a physical layer of host or client link controllers, receivers, or drivers. At least one link controller or driver residing in the host device is coupled to the client device through the communications path or link, and is configured to generate, transmit, and receive packets forming the communications protocol, and to form digital presentation data into one or more types of data packets. The interface provides for bi-directional transfer of information between the host and client, which can reside within a common overall housing or support structure.

[0021] The implementation is generally all digital in nature with the exception of differential drivers and receivers which can be easily implemented on a digital Complementary Metal-Oxide Semiconductor (CMOS) chip, requires a few as 6 signals, and operates at almost any data rate that is convenient for the system designer. The simple physical and link layer protocol makes it easy to integrate, and this simplicity plus a hibernation state enables the portable system to have very low system power consumption.

[0022] To aid in use and acceptance, the interface will add very little to the cost of a device, will allow for consumption of very little power while able to power displays through the interface using standard battery voltages, and can accommodate devices having a pocket-able form-factor. The interface is scalable to support resolutions beyond HDTV, supports simultaneous stereo video and 7.1 audio to a display device, performs conditional updates to any screen region, and supports multiple data types in both directions.

[0023] In further aspects of embodiments of the invention, at least one client link controller, receiver, device, or driver is disposed in the client device and is coupled to the host device through the communications path or link. The client link controller is also configured to generate, transmit, and receive packets forming the communications protocol, and to form digital presentation data into one or more types of data packets. Generally, the host or link controller employs a state machine for processing data packets used in commands or certain types of signal preparation and inquiry processing,

but can use a slower general purpose processor to manipulate data and some of the less complex packets used in the communication protocol. The host controller comprises one or more differential line drivers; while the client receiver comprises one or more differential line receivers coupled to the communication path.

[0024] The packets are grouped together within media frames that are communicated between the host and client devices having a pre-defined fixed length with a pre-determined number of packets having different variable lengths. The packets each comprise a packet length field, one or more packet data fields, and a cyclic redundancy check field. A Sub-frame Header Packet is transferred or positioned at the beginning of transfers of other packets from the host link controller. One or more Video Stream type packets and Audio Stream type packets are used by the communications protocol to transfer video type data and audio type data, respectively, from the host to the client over a forward link for presentation to a client device user. One or more Reverse Link Encapsulation type packets are used by the communications protocol to transfer data from the client device to the host link controller. These transfer in some embodiments include the transfer of data from internal controllers having at least one MDDI device to internal video screens. Other embodiments include transfer to internal sound systems, and transfers from various input devices including joysticks and complex keyboards to internal host devices.

[0025] Filler type packets are generated by the host link controller to occupy periods of forward link transmission that do not have data. A plurality of other packets is used by the communications protocol to transfer video information. Such packets include Color Map, Bit Block Transfer, Bitmap Area Fill, Bitmap Pattern Fill, and Transparent Color Enable type packets. User-Defined Stream type packets are used by the communications protocol to transfer interface-user defined data. Keyboard Data and Pointing Device Data type packets are used by the communications protocol to transfer data to or from user input devices associated with said client device. A Link Shutdown type packet is used by the communications protocol to terminate the transfer of data in either direction over said communication path.

[0026] Display Power State Packets are generated to provide a structure, means, or method for placing specific display controller hardware into a low power state when a client such as a display is not being used or in current active use, in order to minimize the system power consumption or drain on system resources. In one embodiment, a

client indicates an ability to respond to Display Power State Packets using bit 9 of a Client Feature Capability Indicators field of a Client Capability Packet.

[0027] The format of one embodiment for a Display Power State Packet, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Power State, and Cyclic Redundancy Check (CRC) fields. This type of packet is generally identified as a Type 75 packet in the 2-byte type field. The 2-byte hClient ID field contains information or values that are reserved for a Client ID. The Power State field, specifies the information used to place a specific display into the specified power state according to the value of certain pre-selected bits. A 2 byte CRC field specifies or contains the CRC of all bytes in the packet including the Packet Length.

[0028] The communication path generally comprises or employs a cable having a series of four or more conductors and a shield. In addition, printed wires or conductors can be used, as desired, with some residing on flexible substrates.

[0029] The host link controller requests display capabilities information from the client device in order to determine what type of data and data rates said client is capable of accommodating through said interface. The client link controller communicates display or presentation capabilities to the host link controller using at least one Client Capability type packet. Multiple transfer modes are used by the communications protocol with each allowing the transfer of different maximum numbers of bits of data in parallel over a given time period, with each mode selectable by negotiation between the host and client link controllers. These transfer modes are dynamically adjustable during transfer of data, and the same mode need not be used on the reverse link as is used on the forward link.

[0030] In other aspects of some embodiments of the invention, the host device comprises a wireless communications device, such as a wireless telephone, a wireless PDA, or a portable computer having a wireless modem disposed therein. A typical client device comprises a portable video display such as a micro-display device, and/or a portable audio presentation system. Furthermore, the host may use storage means or elements to store presentation or multimedia data to be transferred to be presented to a client device user.

[0031] In still other aspects of some embodiments, the host device comprises a controller or communication link control device with drivers as described below residing within a portable electronic device such as a wireless communications device,

such as a wireless telephone, a wireless PDA, or a portable computer. A typical client device in this configuration comprises a client circuit or integrated circuit or module coupled to the host and residing within the same device, and to an internal video display such as a high resolution screen for a mobile phone, and/or a portable audio presentation system, or in the alternative some type of input system or device.

BRIEF DESCRIPTION OF THE DRAWINGS

- [0032] Further features and advantages, as well as the structure and operation of various embodiments of the invention, are described in detail below with reference to the accompanying drawings. In the drawings, like reference numbers generally indicate identical, functionally similar, and/or structurally similar elements or processing steps.
- [0033] FIG. 1A illustrates a basic environment in which embodiments of the invention might operate including the use of a micro-display device, or a projector, in conjunction with a portable computer or other data processing device.
- [0034] FIG. 1B illustrates a basic environment in which embodiments of the invention might operate including the use of a micro-display device or a projector, and audio presentation elements used in conjunction with a wireless transceiver.
- [0035] FIG. 2A illustrates a basic environment in which embodiments of the invention might operate including the use of internal display or audio presentation devices used in a portable computer.
- [0036] FIG. 2B illustrates a basic environment in which embodiments of the invention might operate including the use of internal display or audio presentation elements used in a wireless transceiver.
- [0037] FIG. 3 illustrates the overall concept of a MDDI with a host and client interconnection.
- [0038] FIG. 4 illustrates the structure of a packet useful for realizing data transfers from a client device to a host device.
- [0039] FIG. 5 illustrates the use of an MDDI link controller and the types of signals passed between a host and a client over the physical data link conductors for a Type 1 interface.
- [0040] FIG. 6 illustrates the use of an MDDI link controller and the types of signals passed between a host and a client over the physical data link conductors for Types 2, 3, and 4 interfaces.

[0041] FIG. 7 illustrates the structure of frames and sub-frames used to implement the interface protocol.

[0042] FIG. 8 illustrates the general structure of packets used to implement the interface protocol.

[0043] FIG. 9 illustrates the format of a Sub-frame Header Packet.

[0044] FIG. 10 illustrates the format and contents of a Filler Packet.

[0045] FIG. 11 illustrates the format of a Video Stream Packet.

[0046] FIGs. 12A-12E illustrate the format and contents for the Video Data Format Descriptor used in FIG. 11.

[0047] FIG. 13 illustrates the use of packed and unpacked formats for data.

[0048] FIG. 14 illustrates the format of an Audio Stream Packet.

[0049] FIG. 15 illustrates the use of byte-aligned and packed PCM formats for data.

[0050] FIG. 16 illustrates the format of a User-Defined Stream Packet.

[0051] FIG. 17 illustrates the format of a Color Map Packet.

[0052] FIG. 18 illustrates the format of a Reverse Link Encapsulation Packet.

[0053] FIG. 19 illustrates the format of a Client Capability Packet.

[0054] FIG. 20 illustrates the format of a Keyboard Data Packet.

[0055] FIG. 21 illustrates the format of a Pointing Device Data Packet.

[0056] FIG. 22 illustrates the format of a Link Shutdown Packet.

[0057] FIG. 23 illustrates the format of a Client Request and Status Packet.

[0058] FIG. 24 illustrates the format of a Bit Block Transfer Packet.

[0059] FIG. 25 illustrates the format of a Bitmap Area Fill Packet.

[0060] FIG. 26 illustrates the format of a Bitmap Pattern Fill Packet.

[0061] FIG. 27 illustrates the format of a Communication Link Data Channel Packet.

[0062] FIG. 28 illustrates the format of a Display Power State Packet.

[0063] FIG. 29 illustrates the format of a Perform Type Handoff Packet.

[0064] FIG. 30 illustrates the format of a Forward Audio Channel Enable Packet.

[0065] FIG. 31 illustrates the format of a Reverse Audio Sample Rate Packet.

[0066] FIG. 32 illustrates the format of a Digital Content Protection Overhead Packet.

[0067] FIG. 33 illustrates the format of a Transparent Color Enable Packet.

[0068] FIG. 34 illustrates the format of a Round Trip Delay Measurement Packet.

[0069] FIG. 35 illustrates the timing of events during the Round Trip Delay Measurement Packet.

- [0070] FIG. 36 illustrates a sample implementation of a CRC generator and checker useful for implementing the invention.
- [0071] FIG. 37A illustrates the timing of CRC signals for the apparatus of FIG. 36 when sending data packets.
- [0072] FIG. 37B illustrates the timing of CRC signals for the apparatus of FIG. 36 when receiving data packets.
- [0073] FIG. 38 illustrates processing steps for a typical service request with no contention.
- [0074] FIG. 39 illustrates processing steps for a typical service request asserted after the link restart sequence has begun, contending with link start.
- [0075] FIG. 40 illustrates how a data sequence can be transmitted using DATA-STB encoding.
- [0076] FIG. 41 illustrates circuitry useful for generating the DATA and STB signals from input data at the host, and then recovering the data at the client.
- [0077] FIG. 42 illustrates drivers and terminating resistors useful for implementing one embodiment.
- [0078] FIGs. 43A-43C illustrate steps and signal levels employed by a client to secure service from the host and by the host to provide such service.
- [0079] FIG. 44 illustrates relative spacing between transitions on the Data0, other Data lines (DataX), and the Strobe lines (Stb).
- [0080] FIG. 45 illustrates the presence of a delay in response that can occur when a host disables the host driver after transferring a packet.
- [0081] FIG. 46 illustrates the presence of a delay in response that can occur when a host enables the host driver to transfer a packet.
- [0082] FIG. 47 illustrates leakage current analysis.
- [0083] FIG. 48 illustrates switching characteristics and relative timing relationships for host and client output enabled and disable time.
- [0084] FIG. 49 illustrates a high level diagram of signal processing steps and conditions by which synchronization can be implemented using a state machine.
- [0085] FIG. 50 illustrates typical amounts of delay encountered for signal processing on the forward and reverse paths in a system employing the MDDI.
- [0086] FIG. 51 illustrates marginal round trip delay measurement.
- [0087] FIG. 52A illustrates Reverse Link data rate changes.

- [0088] FIG. 52B illustrates an example of advanced reverse data sampling.
- [0089] FIG. 53 illustrates a graphical representation of values of the Reverse Rate Divisor versus forward link data rate.
- [0090] FIGs. 54A and 54B illustrate steps undertaken in the operation of an interface.
- [0091] FIG. 55 illustrates an overview of the interface apparatus processing packets.
- [0092] FIG. 56 illustrates the format of a Forward Link Packet.
- [0093] FIG. 57 illustrates typical values for propagation delay and skew in an Type 1 Link interface.
- [0094] FIG. 58 illustrates Data, Stb, and Clock Recovery Timing on a Type 1 Link for exemplary signal processing through the interface.
- [0095] FIG. 59 illustrates typical values for propagation delay and skew in Type 2, Type 3 or Type 4 Link interfaces.
- [0096] FIGs. 60A, 60B, and 60C illustrate different possibilities for the timing of two data signals and MDDI_Stb with respect to each other, being ideal, early, and late, respectively.
- [0097] FIG. 61 illustrates interface pin assignments exemplary connectors used with a Type-1/Type 2 interfaces.
- [0098] FIGs. 62A and 62B illustrate possible MDDI_Data and MDDI_Stb waveforms for both Type 1 and Type 2 Interfaces, respectively.
- [0099] FIG. 63 illustrates a high level diagram of alternative signal processing steps and conditions by which synchronization can be implemented using a state machine.
- [0100] FIG. 64 illustrates exemplary relative timing between a series of clock cycles and the timing of a various reverse link packets bits and divisor values.
- [0101] FIG. 65 illustrates exemplary error code transfer processing.
- [0102] FIG. 66 illustrates apparatus useful for error code transfer processing.
- [0103] FIG. 67A illustrates error code transfer processing for code overloading.
- [0104] FIG. 67B illustrates error code transfer processing for code reception.
- [0105] FIG. 68A illustrates processing steps for a host initiated wake-up.
- [0106] FIG. 68B illustrates processing steps for a client initiated wake-up.
- [0107] FIG. 68C illustrates processing steps for host and client initiated wake-up with contention.
- [0108] FIG. 69 illustrates the format of a Request Virtual Control Panel (VCP) Feature Packet.

- [00109] FIG. 70 illustrates the format of a VCP Feature Reply Packet.
- [00110] FIG. 71 illustrates the format of a VCP Feature Reply List.
- [00111] FIG. 72 illustrates the format of a Set VCP Feature Packet.
- [00112] FIG. 73 illustrates the format of a Request Valid Parameter Packet.
- [00113] FIG. 74 illustrates the format of a Valid Parameter Reply Packet.
- [00114] FIG. 75 illustrates the format of a Scaled Video Stream Capability Packet.
- [00115] FIG. 76 illustrates the format of a Scaled Video Stream Setup Packet.
- [00116] FIG. 77 illustrates the format of a Scaled Video Stream Acknowledgement Packet.
- [00117] FIG. 78 illustrates the format of a Scaled Video Stream Packet.
- [00118] FIG. 79 illustrates the format of a Request Specific Status Packet.
- [00119] FIG. 80 illustrates the format of a Valid Status Reply List Packet.
- [00120] FIG. 81 illustrates the format of a Personal Display Capability Packet.
- [00121] FIG. 82 illustrates elements in the Points of Field Curvature List.
- [00122] FIG. 83 illustrates the format of a Client Error Report Packet.
- [00123] FIG. 84 illustrates the format of an Error Report List item.
- [00124] FIG. 85 illustrates the format of a Client Identification Packet.
- [00125] FIG. 86 illustrates the format of an Alternate Display Capability Packet.
- [00126] FIG. 87 illustrates the format of a Register Access Packet.
- [00127] FIG. 88A-88C illustrate use of two display buffers to reduce visible artifacts.
- [00128] FIG. 89 illustrates two buffers with display refresh faster than image transfer.
- [00129] FIG. 90 illustrates two buffers with display refresh slower than image transfer.
- [00130] FIG. 91 illustrates two buffers with display refresh much faster than image transfer.
- [00131] FIG. 92 illustrates three buffers with display refresh faster than image transfer.
- [00132] FIG. 93 illustrates three buffers with display refresh slower than image transfer.
- [00133] FIG. 94 illustrates one buffer with display refresh faster than image transfer.
- [00134] FIG. 95 illustrates host-client connection via daisy-chain and hub.
- [00135] FIG. 96 illustrates client devices connected via a combination of hubs and daisy chains.
- [00136] FIG. 97 illustrates a color map.

DETAILED DESCRIPTION

I. Overview

[00137] A general intent of the invention is to provide a MDDI, as discussed below, which results in or provides a cost-effective, low power consumption, transfer mechanism that enables high- or very-high- speed data transfer over a short-range communication link between a host device and a client device, such as a display element, using a "serial" type of data link or channel. This mechanism lends itself to implementation with miniature connectors and thin flexible cables which are especially useful in connecting internal (interior to a housing or support frame) display or output elements or devices, or input devices to a central controller or communication element or device. In addition, this connection mechanism is very useful for connecting external display elements or devices such as wearable micro-displays (goggles or projectors) or other types of visual, audible, tactile information presentation devices to portable computers, wireless communication devices, or entertainment devices.

[00138] Although the terms Mobile and Display are associated with the naming of the protocol, it is to be understood that this is for convenience only in terms of having a standard name easily understood by those skilled in the art working with the interface and protocol. As it will relate to a Video Electronics Standards Association (VESA) standard and various applications of that standard. However, it will be readily understood after a review of the embodiments presented below that many non-mobile and non-display related applications will benefit from application of this protocol, resulting interface structure, or transfer mechanism, and the MDDI label is not intended to imply any limitations to the nature or usefulness of the invention or its various embodiments.

[00139] An advantage of embodiments of the invention is that a technique is provided for data transfer that is low in complexity, low cost, has high reliability, fits well within the environment of use, and is very robust, while remaining very flexible.

[00140] Embodiments of the invention can be used in a variety of situations to communicate or transfer large quantities of data, generally for audio, video, or multimedia applications from a host or source device where such data is generated, manipulated, such as for transfer to specific devices, or otherwise processed, or stored; to a client or receiving device, such as a video display or projection element, audio

speakers, or other presentation device at a high rate. A typical application, which is discussed below, is the transfer of data from either a portable computer or a wireless telephone or modem to a visual display device such as a small video screen or a wearable micro-display appliance, such as in the form of goggles or helmets containing small projection lenses and screens, or from a host to client device within such components. That is, from a processor or controller to an internal screen or other presentation element, as well as from various internal, or external input devices employing a client to an internally located (collocated within same device housing or support structure) host, or connected thereto by a cable or conductors.

[00141] The characteristics or attributes of the MDDI are such that they are independent of specific display or presentation technology. This is a highly flexible mechanism for transferring data at a high rate without regards to the internal structure of that data nor the functional aspects of the data or commands it implements. This allows the timing of data packets being transferred to be adjusted to adapt to the idiosyncrasies of particular client devices, such as for unique display desires for certain devices, or to meet the requirements of combined audio and video for some Audio Visual (A-V) systems, or for certain input devices such as joysticks, touch pads, and so forth. The interface is very display element or client device agnostic, as long as the selected protocol is followed. In addition, the aggregate serial link data, or data rate, can vary over several orders of magnitude which allows a communication system or host device designer to optimize the cost, power requirements, client device complexity, and client device update rates.

[00142] The data interface is presented primarily for use in transferring large amounts of high rate data over a "wired" signal link or small cable. However, some applications may take advantage of a wireless link as well, including optical based links, provided it is configured to use the same packet and data structures developed for the interface protocol, and can sustain the desired level of transfer at low enough power consumption or complexity to remain practical.

II. Environment

[00143] A typical application can be seen in FIGs. 1A and 1B where a portable or laptop computer 100 and wireless telephone or PDA device 102 are shown communicating data with display devices 104 and 106, respectively, along with audio reproduction

systems 108 and 112. In addition, FIG. 1A shows potential connections to a larger display or screen 114 or an image projector 116, which are only shown in one figure for clarity, but are connectable to wireless device 102 as well. The wireless device can be currently receiving data or have previously stored a certain amount of multimedia type data in a memory element or device for later presentation for viewing and/or hearing by an end user of the wireless device. Since a typical wireless device is used for voice and simple text communications most of the time, it has a rather small display screen and simple audio system (speakers) for communicating information to the device 102 user.

[00144] Computer 100 has a much larger screen, but still inadequate external sound system, and still falls short of other multimedia presentation devices such as a high definition television, or movie screens. Computer 100 is used for purposes of illustration and other types of processors, interactive video games, or consumer electronics devices can also be used with the invention. Computer 100 can employ, but is not limited to or by, a wireless modem or other built in device for wireless communications, or be connected to such devices using a cable or wireless link, as desired.

[00145] This makes presentation of more complex or "rich" data a less than a useful or enjoyable experience. Therefore, the industry is developing other mechanisms and devices to present the information to end users and provide a minimum level of desired enjoyment or positive experience.

[00146] As previously discussed above, several types of display devices have or are currently being developed for presenting information to end users of device 100. For example, one or more companies have developed sets of wearable goggles that project an image in front of the eyes of a device user to present a visual display. When correctly positioned such devices effectively "project" a virtual image, as perceived by a user's eyes, which is much larger than the element providing the visual output. That is, a very small projection element allows the eye(s) of the user to "see" images on a much larger scale than possible with typical Liquid Crystal Display (LCD) screens and the like. The use of larger virtual screen images also allows the use of much higher resolution images than possible with more limited LCD screen displays. Other display devices could include, but are not limited to, small LCD screens or various flat panel display elements, projection lenses and display drivers for projecting images on a surface, and so forth.

[00147] There may also be additional elements connected to or associated with the use of wireless device 102 or computer 100 for presenting an output to another user, or to another device which in turn transfers the signals elsewhere or stores them. For example, data may be stored in flash memory, in optical form, for example using a writeable CD media or on magnetic media such as in a magnetic tape recorder and similar devices, for later use.

[00148] In addition, many wireless devices and computers now have built-in MP3 music decoding capabilities, as well as other advanced sound decoders and systems. Portable computers utilize CD and DVD playback capabilities as a general rule, and some have small dedicated flash memory readers for receiving pre-recorded audio files. The issue with having such capabilities is that digital music files promise a highly increased feature rich experience, but only if the decoding and playback process can keep pace. The same holds true for the digital video files.

[00149] To assist with sound reproduction, external speakers 114 are shown in FIG. 1A, which could also be accompanied by addition elements such as sub-woofers, or "surround-sound" speakers for front and rear sound projection. At the same time, speakers or earphones 108 are indicated as built-in to the support frame or mechanism of micro-display device 106 of FIG. 1B. As would be known, other audio or sound reproduction elements can be used including power amplification or sound shaping devices.

[00150] In any case, as discussed above, when one desires to transfer high quality or high resolution image data and high quality audio information or data signals from a data source to an end user over one or more communication links 110, a high data rate is required. That is, transfer link 110 is clearly a potential bottleneck in the communication of data as discussed earlier, and is limiting system performance, since current transfer mechanisms do not achieve the high data rates typically desired. As discussed above for example, for higher image resolutions such as 1024 by 1024 pixels, with color depths of 24-32 bits per pixel and at data rates of 30 fps, the data rates can approach rates in excess of 755 Mbps or more. In addition, such images may be presented as part of a multimedia presentation which includes audio data and potentially additional signals dealing with interactive gaming or communications, or various commands, controls, or signals, further increasing the quantity or data and the data rate.

- [00151] It is also clear that fewer cables or interconnections required for establishing a data link, means that mobile devices associated with a display are easier to use, and more likely to be adopted by a larger user base. This is especially true where multiple devices are commonly used to establish a full audio-visual experience, and more especially as the quality level of the displays and audio output devices increases.
- [00152] Another typical application related to many of the above and other improvements in video screens and other output or input devices can be seen in FIGs. 1C and 1D where a portable or laptop computer 130 and wireless telephone or PDA device 140 are shown communicating data with "internal" display devices 134 and 144, respectively, along with audio reproduction systems 136 and 146.
- [00153] In FIGs. 2A and 2B, small cut-away sections of the overall electronic devices or products are used to show the location of one or more internal hosts and controllers in one portion of the device with a generalized communication link, here 138 and 148, respectively, connecting them to the video display elements or screens having the corresponding clients, across a rotating joint of some known type used throughout the electronics industry today. One can see that the amount of data involved in these transfers requires a large number of conductors to comprise links 138 and 148. It is estimated that such communication links are approaching 90 or more conductors in order to satisfy today's growing needs for utilizing advanced color and graphical interfaces, display elements, on such devices because of the types of parallel or other known interface techniques available for transferring such data.
- [00154] Unfortunately, the higher data rates exceed current technology available for transferring data. Both in terms of the raw amount of data needing to be transferred per unit time, and in terms of manufacturing reliable cost effective physical transfer mechanisms.
- [00155] What is needed is a technique, a structure, means or method, for transferring data at higher rates for the data transfer link or communication path between presentation elements and the data source, which allows for consistently low(er) power, light weight, and as simple and economical a cabling structure as possible. Applicants have developed a new technique, or method and apparatus, to achieve these and other goals to allow an array of mobile, portable, or even fixed location devices to transfer data to desired displays, micro-displays, or audio transfer elements, at very high data rates, while maintaining a desired low power consumption, and complexity.

III. High Rate Digital Data Interface System Architecture

[00156] In order to create and efficiently utilize a new device interface, a signal protocol and system architecture has been formulated that provides a very high data transfer rate using low power signals. The protocol is based on a packet and common frame structure, or structures linked together to form a protocol for communicating a pre-selected set of data or data types along with a command or operational structure imposed on the interface.

A. Overview

[00157] The devices connected by or communicating over the MDDI link are called the host and client, with the client typically being a display device of some type, although other output and input devices are contemplated. Data from the host to the display travels in the forward direction (referred to as forward traffic or link), and data from the client to the host travels in the reverse direction (reverse traffic or link), as enabled by the host. This is illustrated in the basic configuration shown in FIG. 3. In FIG. 3, a host 202 is connected to a client 204 using a bi-directional communication channel 206 which is illustrated as comprising a forward link 208 and a reverse link 210. However, these channels are formed by a common set of conductors whose data transfer is effectively switched between the forward or reverse link operations. This allows for greatly reduced numbers of conductors, immediately addressing one of the many problems faced with current approaches to high speed data transfer in low power environments such as for mobile electronic devices.

[00158] As discussed elsewhere, the host comprises one of several types of devices that can benefit from using the present invention. For example, host 202 could be a portable computer in the form of a handheld, laptop, or similar mobile computing device. It could also be a PDA, a paging device, or one of many wireless telephones or modems. Alternatively, host 202 could be a portable entertainment or presentation device such as a portable DVD or CD player, or a game playing device.

[00159] Furthermore, the host can reside as a host device or control element in a variety of other widely used or planned commercial products for which a high speed communication link is desired with a client. For example, a host could be used to

transfer data at high rates from a video recording device to a storage based client for improved response, or to a high resolution larger screen for presentations. An appliance such as a refrigerator that incorporates an onboard inventory or computing system and/or Bluetooth connections to other household devices, can have improved display capabilities when operating in an internet or Bluetooth connected mode, or have reduced wiring needs for in-the-door displays (a client) and keypads or scanners (client) while the electronic computer or control systems (host) reside elsewhere in the cabinet. In general, those skilled in the art will appreciate the wide variety of modern electronic devices and appliances that may benefit from the use of this interface, as well as the ability to retrofit older devices with higher data rate transport of information utilizing limited numbers of conductors available in either newly added or existing connectors or cables.

[00160] At the same time, client 204 could comprise a variety of devices useful for presenting information to an end user, or presenting information from a user to the host. For example, a micro-display incorporated in goggles or glasses, a projection device built into a hat or helmet, a small screen or even holographic element built into a vehicle, such as in a window or windshield, or various speaker, headphone, or sound systems for presenting high quality sound or music. Other presentation devices include projectors or projection devices used to present information for meetings, or for movies and television images. Another example would be the use of touch pads or sensitive devices, voice recognition input devices, security scanners, and so forth that may be called upon to transfer a significant amount of information from a device or system user with little actual "input" other than touch or sound from the user. In addition, docking stations for computers and car kits or desk-top kits and holders for wireless telephones may act as interface devices to end users or to other devices and equipment, and employ either clients (output or input devices such as mice) or hosts to assist in the transfer of data, especially where high speed networks are involved.

[00161] However, those skilled in the art will readily recognize that the present invention is not limited to these devices, there being many other devices on the market, and proposed for use, that are intended to provide end users with high quality images and sound, either in terms of storage and transport or in terms of presentation at playback. The present invention is useful in increasing the data throughput between various

elements or devices to accommodate the high data rates needed for realizing the desired user experience.

[00162] The inventive MDDI and communication signal protocol may be used to simplify the interconnect between a host processor, controller, or circuit component (for example), and a display within a device or device housing or structure (referred to as an internal mode) in order to reduce the cost or complexity and associated power and control requirements or constraints of these connections, and to improve reliability, not just for connection to or for external elements, devices, or equipment (referred to as an external mode).

[00163] The aggregate serial link data rate on each signal pair used by this interface structure can vary over many orders of magnitude, which allows a system or device designer to easily optimize cost, power, implementation complexity, and the display update rate for a given application or purpose. The attributes of MDDI are independent of display or other presentation device (target client) technology. The timing of data packets transferred through the interface can be easily adjusted to adapt to idiosyncrasies of particular clients such as display devices, sound systems, memory and control elements, or combined timing requirements of audio-video systems. While this makes it possible to have a system with a very small power consumption, it is not a requirement of the various clients to have frame buffers in order to make use of the MDDI protocol at least at some level.

B. Interface Types

[00164] The MDDI is contemplated as addressing at least four, and potentially more, somewhat distinct physical types of interfaces found in the communications and computer industries. These are labeled simply as Type 1, Type 2, Type 3, and Type 4, although other labels or designations may be applied by those skilled in the art depending upon the specific applications they are used for or industry they are associated with. For example, simple audio systems use fewer connections than more complex multimedia systems, and may reference features such as "channels" differently, and so forth.

[00165] The Type 1 interface is configured as a 6-wire, or other type of conductor or conductive element, interface which makes it suitable for mobile or wireless telephones, PDAs, electronic games, and portable media players, such as CD players, or MP3

players, and similar devices or devices used on similar types of electronic consumer technology. In one embodiment, a an interface can be configured as an 8-wire (conductor) interface which is more suitable for laptop, notebook, or desktop personal computers and similar devices or applications, not requiring rapid data updates and which do not have a built-in MDDI link controller. This interface type is also distinguishable by the use of an additional two-wire Universal Serial Bus (USB) interface, which is extremely useful in accommodating existing operating systems or software support found on most personal computers.

[00166] Type 2, Type 3, and Type 4 interfaces are suitable for high performance clients or devices and use larger more complex cabling with additional twisted-pair type conductors to provide the appropriate shielding and low loss transfers for data signals.

[00167] The Type 1 interface passes signals which can comprise display, audio, control, and limited signaling information, and is typically used for mobile clients or client devices that do not require high-resolution full-rate video data. A Type 1 interface can easily support Super Video Graphics Array (SVGA) resolution at 30 fps plus 5.1 channel audio, and in a minimum configuration might use only three wire pairs total, two pairs for data transmission and one pair for power transfer. This type of interface is primarily intended for devices, such as mobile wireless devices, where a USB host is typically not available within the such device for connection and transfer of signals. In this configuration, the mobile wireless device is a MDDI host device, and acts as the "master" that controls the communication link from the host, which generally sends data to the client (forward traffic or link) for presentation, display or playback.

[00168] In this interface, a host enables receipt of communication data at the host from the client (reverse traffic or link) by sending a special command or packet type to the client that allows it to take over the bus (link) for a specified duration and send data to the host as reverse packets. This is illustrated in FIG. 4, where a type of packet referred to as an encapsulation packet (discussed below) is used to accommodate the transfer of reverse packets over the transfer link, creating the reverse link. The time interval allocated for the host to poll the client for data is pre-determined by the host, and is based on the requirements of each specified application. This type of half-duplex bi-directional data transfer is especially advantageous where a USB port is not available for transfer of information or data from the client.

[00169] High-performance displays capable of HDTV type or similar high resolutions require around 1.5 Gbps rate data streams in order to support full-motion video. The Type 2 interface supports high data rates by transmitting 2 bits in parallel, the Type 3 by transmitting 4 bits in parallel, and the Type 4 interface transfers 8 bits in parallel. Type 2 and Type 3 interfaces use the same cable and connector as Type 1 but can operate at twice and four times the data rate to support higher-performance video applications on portable devices. A Type 4 interface is suited for very high performance clients or displays and requires a slightly larger cable that contains additional twisted-pair data signals.

[00170] The protocol used by the MDDI allows each Type 1, 2, 3, or 4 host to generally communicate with any Type 1, 2, 3, or 4 client by negotiating what is the highest data rate possible that can be used. The capabilities or available features of what can be referred to as the least capable device is used to set the performance of the link. As a rule, even for systems where the host and client are both capable using Type 2, Type 3, or Type 4 interfaces, both begin operation as a Type 1 interface. The host then determines the capability of the target client, and negotiates a hand-off or reconfiguration operation to either Type 2, Type 3, or Type 4 mode, as appropriate for the particular application.

[00171] It is generally possible for the host to use the proper link-layer protocol (discussed further below) and step down or again reconfigure operation at generally any time to a slower mode to save power or to step up to a faster mode to support higher speed transfers, such as for higher resolution display content. For example, a host may change interface types when the system switches from a power source such as a battery to AC power, or when the source of the display media switches to a lower or higher resolution format, or a combination of these or other conditions or events may be considered as a basis for changing an interface type, or transfer mode.

[00172] It is also possible for a system to communicate data using one mode in one direction and another mode in another direction. For example, a Type 4 interface mode could be used to transfer data to a display at a high rate, while a Type 1 mode is used when transferring data to a host device from peripheral devices such as a keyboard or a pointing device. It will be appreciated by one skilled in the art that hosts and clients may communicate outgoing data at different rates.

[00173] Often, users of the MDDI protocol may distinguish between an “external” mode and an “internal” mode. An external mode describes the use of the protocol and interface to connect a host in one device to a client outside of that device that is up to about 2 meters or so from the host. In this situation, the host may also send power to the external client so that both devices can easily operate in a mobile environment. An internal mode describes when the host is connected to a client contained inside the same device, such as within a common housing or support frame or structure of some kind. An example would be applications within a wireless phone or other wireless device, or a portable computer or gaming device where the client is a display or display driver, or an input device such as a keypad or touch-pad, or a sound system, and the host is a central controller, graphics engine, or CPU element. Since a client is located much closer to the host in internal mode applications as opposed to external mode applications, there are generally no requirements discussed for the power connection to the client in such configurations.

C. Physical Interface Structure

[00174] The general disposition of a device or link controller for establishing communications between host and client devices is shown in FIGs. 5 and 6. In FIGs. 5 and 6, a MDDI link controller 402 and 502 is shown installed in a host device 202 and a MDDI link controller 404 and 504 is shown installed in a client device 204. As before, host 202 is connected to a client 204 using a bi-directional communication channel 406 comprising a series of conductors. As discussed below, both the host and client link controllers can be manufactured as an integrated circuit using a single circuit design that can be set, adjusted, or programmed to respond as either a host controller (driver) or a client controller (receiver). This provides for lower costs due to larger scale manufacturing of a single circuit device.

[00175] In FIG. 6, a MDDI link controller 502 is shown installed in a host device 202' and a MDDI link controller 504 is shown installed in a client device 204'. As before, host 202' is connected to a client 204' using a bi-directional communication channel 506 comprising a series of conductors. As discussed before, both the host and client link controllers can be manufactured using a single circuit design.

[00176] Signals passed between a host and a client, such as a display device, over the MDDI link, or the physical conductors used, are also illustrated in FIGs. 5 and 6. As

seen in FIGs. 5 and 6, the primary path or mechanism for transferring data through the MDDI uses data signals labeled as MDDI_Data0+/- and MDDI_Stb+/- . Each of these is low voltage data signals that are transferred over a differential pair of wires in a cable. There is only one transition on either the MDDI_Data0 pair or the MDDI_Stb pair for each bit sent over the interface. This is a voltage based transfer mechanism not current based, so static current consumption is nearly zero. The host drives the MDDI_Stb signals to the client display.

[00177] While data can flow in both the forward and reverse directions over the MDDI_Data pairs, that is, it is a bi-directional transfer path, the host is the master or controller of the data link. The MDDI_Data0 and MDDI_Stb signal paths are operated in a differential mode to maximize noise immunity. The data rate for signals on these lines is determined by the rate of the clock sent by the host, and is variable over a range of about 1 kbps up to 400 Mbps or more.

[00178] The Type 2 interface contains one additional data pair or conductors or paths beyond that of the Type 1, referred to as MDDI_Data1+/- . The Type 3 interface contains two additional data pairs or signal paths beyond that of the Type 2 interface referred to as MDDI_Data2+/-, and MDDI_Data3+/- . The Type 4 interface contains four more data pairs or signal paths beyond that of the Type 3 interface referred to as: MDDI_data4+/-, MDDI_Data5+/-, MDDI_Data6+/-, and MDDI_Data7+/-, respectively. In each of the above interface configurations, a host can send power to the client or display using the wire-pair or signals designated as HOST_Pwr and HOST_Gnd. As discussed further below, power transfer can also be accommodated, if desired, in some configurations on the MDDI_data4+/-, MDDI_Data5+/-, MDDI_Data6+/-, or MDDI_Data7+/- conductors when an interface "Type" is being used that employs fewer conductors than are available or present for the other modes. This Power transfer is generally employed for external modes, there generally being no need for internal modes, although some applications may differ.

[00179] A summary of the signals passed between the host and client (display) over the MDDI link for various modes are illustrated in Table I, below, in accordance with the interface type.

Table I

Type 1	Type 2	Type 3	Type 4
HOST_Pwr/Gnd	HOST_Pwr/Gnd	HOST_Pwr/Gnd	HOST_Pwr/Gnd
MDDI_Stb+/-	MDDI_Stb+/-	MDDI_Stb+/-	MDDI_Stb+/-
MDDI_Data0+/-	MDDI_Data0+/-	MDDI_Data0+/-	MDDI_Data0+/-
	MDDI_Data1+/-	MDDI_Data1+/-	MDDI_Data1+/-
		MDDI_Data2+/-	MDDI_Data2+/-
		MDDI_Data3+/-	MDDI_Data3+/-
Optional Pwr	Optional Pwr	Optional Pwr	MDDI_Data4+/-
Optional Pwr	Optional Pwr	Optional Pwr	MDDI_Data5+/-
Optional Pwr	Optional Pwr	Optional Pwr	MDDI_Data6+/-
Optional Pwr	Optional Pwr	Optional Pwr	MDDI_Data7+/-

[00180] Also note that the HOST_Pwr/Gnd connections for transfer from the host are provided generally for external modes. Internal applications or modes of operation generally have clients that draw power directly from other internal resources, and do not use MDDI to control power distribution, as would be apparent to one skilled in the art, so such distribution is not discussed in further detail here. However, it is certainly possible to allow power to be distributed through the MDDI to allow for certain kinds of power control, synchronization, or interconnection convenience, for example, as would be understood by one skilled in the art.

Cabling generally used to implement the above structure and operation is nominally on the order of 1.5 meters in length, generally 2 meters or less, and contains three twisted pairs of conductors, each in turn being multi-strand wire, nominally between 32 American Wire Gauge (AWG) to 28 AWG. Although, the wire size is not constrained to this range, as those skilled in the art will appreciate, electrical specifications or constraints should be met for maximum total end-to-end resistance, maximum capacitance per foot, impedance of each pair, and crosstalk.

[00181] A foil shield covering is wrapped or otherwise formed above the entire, here three, set or group of twisted pairs, and a drain wire as an additional drain wire. The twisted pairs and shield drain conductor terminate in the client connector with the shield connected to the shield for the client, and there is an insulating layer, covering the entire cable, as would be well known in the art. The wires are paired as: HOST_Gnd with HOST_Pwr; MDDI_Stb+ with MDDI_Stb-; MDDI_Data0+ with MDDI_Data0-; MDDI_Data1+ with MDDI_Data1-; and so forth. However, a variety of conductors and

cabling can be used, as would be understood in the art, to implement the embodiments of the invention, depending upon specific applications. For example, heavier outside coatings or even metallic layers may be used to protect the cable in some applications, while thinner, flatter conductive ribbon type structures may be well suited to other applications.

D. Data Types and Rates

[00182] To achieve a useful interface for a full range of user experiences and applications, the MDDI provides support for a variety of clients and display information, audio transducers, keyboards, pointing devices, and many other input or output devices that might be integrated into or working in concert with a mobile communication, computation, or display device, along with control information, and combinations thereof. The MDDI is designed to be able to accommodate a variety of potential types of streams of data traversing between the host and client in either the forward or reverse link directions using a minimum number of cables or conductors. Both isochronous streams and asynchronous stream (updates) are supported. Many combinations of data types are possible as long as the aggregate data rate is less than or equal to the maximum desired MDDI link rate, which is limited by the maximum serial rate and number of data pairs employed. These could include, but are not limited to, those items listed in Tables II and III below.

Table II

Transferring from Host to Client		
isochronous video data	720x480, 12bit, 30f/s	~124.5 Mbps
isochronous stereo audio data	44.1kHz, 16bit, stereo	~ 1.4 Mbps
asynchronous graphics data	800x600, 12bit, 10f/s, stereo	~115.2 Mbps
asynchronous control	Minimum	<< 1.0 Mbps

Table III

Transferring from Client to Host		
isochronous voice data	8 kHz, 8bit	<< 1.0 Mbps
isochronous video data	640x480, 12bit, 24f/s	~ 88.5 Mbps
asynchronous status, user input, etc.	minimum	<< 1.0 Mbps

[00183] The interface is not fixed but extensible so that it can support the transfer of a variety of information "types" which includes user-defined data, for future system flexibility. Specific examples of data to be accommodated are: full-motion video, either in the form of full or partial screen bitmap fields or compressed video; static bitmaps at low rates to conserve power and reduce implementation costs; PCM or compressed audio data at a variety of resolutions or rates; pointing device position and selection, and user-definable data for capabilities yet to be defined. Such data may also be transferred along with control or status information to detect device capability or set operating parameters.

[00184] Embodiments of the invention advance the art for use in data transfers that include, but are not limited to: watching a movie (video display and audio); using a personal computer with limited personal viewing (graphics display, sometimes combined with video and audio); playing a video game on a PC, console, or personal device (motion graphics display, or synthetic video and audio); "surfing" the Internet, using devices in the form of a video phone (bi-directional low-rate video and audio), a camera for still digital pictures, or a camcorder for capturing digital video images; using a phone, computer, or PDA docked with a projector to give a presentation or docked with a desktop docking station connected to a video monitor, keyboard, and mouse; and for productivity enhancement or entertainment use with cell phones, smart phones, or PDAs, including wireless pointing devices and keyboard data.

[00185] The high speed data interface as discussed below is presented in terms of providing large amounts of A-V type data over a communication or transfer link which is generally configured as a wire-line or cable type link. However, it will be readily apparent that the signal structure, protocols, timing, or transfer mechanism could be adjusted to provide a link in the form of an optical or wireless media, if it can sustain the desired level of data transfer.

[00186] The MDDI signals use a concept known as the Common Frame Rate (CFR) for the basic signal protocol or structure. The idea behind using of a Common Frame Rate is to provide a synchronization pulse for simultaneous isochronous data streams by sending Sub-frame Header Packets at a rate that can be used to derive clocks or frame timing for multiple streams. The rate at which Sub-frame Header Packets are sent is the Common Frame Rate. A client device can use this Common Frame Rate as a time reference. A low CFR increases channel efficiency by decreasing overhead to transmit

the sub-frame header. On the other hand, a high CFR decreases the latency, and allows a smaller elastic data buffer for audio samples. The CFR of the present inventive interface is dynamically programmable and may be set at one of many values that are appropriate for the isochronous streams used in a particular application. That is, the CF value is selected to best suit the given client and host configuration, as desired.

[00187] The number of bytes generally required per sub-frame, which is adjustable or programmable, for isochronous data streams that are most likely to be used with an application, such as for a video or micro-display are shown in Table IV.

Table IV

Common Frame Rate (CFR) = 300 Hz							
	X	Y	Bit	Frame Rate	Channel	Rate (Mbps)	Bytes/ Sub-frame
Computer Game	720	480	24	30	1	248.832	103680
Computer Graphics	800	600	24	10	1	115.200	48000
Video	640	480	12	29.97 or 30	1	221.184	92160
CD Audio	1	1	16	44100	2	1.4112	588
Voice	1	1	8	8000	1	0.064	26-2/3

[00188] Fractional counts of bytes per sub-frame are easily obtained using a simple programmable M/N counter structure. For example, a count of 26-2/3 bytes per sub-frame is implemented by transferring 2 sub-frames containing 27 bytes each followed by one sub-frame containing 26 bytes. A smaller CFR may be selected to produce an integer number of bytes per sub-frame. However, generally speaking, to implement a simple M/N counter in hardware should require less area within an integrated circuit chip or electronic module used to implement part or all of embodiments of the invention than the area needed for a larger audio sample FIFO buffer.

[00189] An exemplary application that illustrates the impact of different data transfer rates and data types is a Karaoke system. For Karaoke, a system where an end user, or users, sings along with a music video program. Lyrics of the song are displayed somewhere on, typically at the bottom of, a screen so the user knows the words to be sung, and roughly the timing of the song. This application requires a video display with

infrequent graphics updates, and mixing of the user's voice, or voices, with a stereo audio stream.

[00190] If one assumes a common frame rate of 300 Hz, then each sub-frame will consist of 92,160 bytes of video content and 588 bytes of audio content (based on 147 16-bit samples, in stereo) over the forward link to the client, and an average of 26.67 (26-2/3) bytes of voice are sent back from a microphone to the mobile Karaoke machine. Asynchronous packets are sent between the host and the client, possibly a head mounted display. This includes the bottom quarter-screen height being updated with lyric text at 1/30th of second intervals or periods, and other miscellaneous control and status commands sent in sub-frames when lyric text is not being updated.

[00191] Table V, shows how data is allocated within a sub-frame for the Karaoke example. The total rate being used is selected to be about 279 Mbps. A slightly higher rate of 280 Mbps allows about another 400 bytes of data per sub-frame to be transferred which allows the use of occasional control and status messages.

Table V

Element Rate	Overhead Bytes per sub-frame	Media Bytes per sub-frame
Music Video at 640 x 480 pixels and 30 fps	$2 * 28 = 56$	92160
Lyric Text at 640 x 120 pixels and 1 fps Updated in 10 sub-frames, 1/30 sec.	28	23040
CD Audio at 44,100 sps, stereo, 16-bit	$2 * 16 = 32$	588
Voice at 8,000 sps, mono, 8-bit	$28+8+8+(4*16)+(3*27) = 125$	26.67 (27 max)
Sub-frame Header	22	
Total Bytes/CF	263	115815
Total Rate (Mbps)	$(263+115815)*8*300 = 278.5872$	

III.(Continued) High Rate Digital Data Interface System Architecture

E. Link Layer

[00192] Data transferred using the MDDI high-speed serial data signals consists of a stream of time-multiplexed packets that are linked one after the other. Even when a transmitting device has no data to send, a MDDI link controller generally automatically sends filler packets, thus, maintaining a stream of packets. The use of a simple packet

structure ensures reliable isochronous timing for video and audio signals or data streams.

[00193] Groups of packets are contained within signal elements or structures referred to as sub-frames, and groups of sub-frames are contained within signal elements or structures referred to as a media-frame. A sub-frame contains one or more packets, depending on their respective size and data transfer uses, and a media-frame contains one more sub-frames. The largest sub-frame provided by the protocol employed by the embodiments presented here is on the order of $2^{32}-1$ or 4,294,967,295 bytes, and the largest media-frame size then becomes on the order of $2^{16}-1$ or 65,535 sub-frames.

[00194] A special sub-frame header packet contains a unique identifier that appears at the beginning of each sub-frame, as is discussed below. That identifier is also used for acquiring the frame timing at the client device when communication between the host and client is initiated. Link timing acquisition is discussed in more detail below.

[00195] Typically, a display screen is updated once per media-frame when full-motion video is being displayed. The display frame rate is the same as the media-frame rate. The link protocol supports full-motion video over an entire display, or just a small region of full-motion video content surrounded by a static image, depending on the desired application. In some low-power mobile applications, such as viewing web pages or email, the display screen may only need to be updated occasionally. In those situations, it is advantageous to transmit a single sub-frame and then shut down or inactivate the link to minimize power consumption. The interface also supports effects such as stereo vision, and handles graphics primitives.

[00196] Sub-frames allow a system to enable the transmission of high-priority packets on a periodic basis. This allows simultaneous isochronous streams to co-exist with a minimal amount of data buffering. This is one advantage embodiments provide to the display process, allowing multiple data streams (high speed communication of video, voice, control, status, pointing device data, etc.) to essentially share a common channel. It transfers information using relatively few signals. It also enables display-technology-specific actions to exist, such as horizontal sync pulses and blanking intervals for a CRT monitor, or for other client-technology-specific actions.

F. Link Controller

[00197] The MDDI link controller shown in FIGs. 5 and 6 is manufactured or assembled to be a completely digital implementation with the exception of the differential line receivers which are used to receive MDDI data and strobe signals. However, even the differential line drivers and receivers can be implemented in the same digital integrated circuits with the link controller, such as when making a CMOS type IC. No analog functions or Phase Lock Loops (PLLs) are required for bit recovery or to implement the hardware for the link controller. The host and client link controllers contain very similar functions, with the exception of the client interface which contains a state machine for link synchronization. Therefore, the embodiments of the invention allow the practical advantage of being able to create a single controller design or circuit that can be configured as either a host or client, which can reduce manufacturing costs for the link controllers, as a whole.

IV. Interface Link Protocol

A. Frame structure

[00198] The signal protocol or frame structure used to implement the forward link communication for packet transfer is illustrated in FIG. 7. As shown in FIG. 7, information or digital data is grouped into elements known as packets. Multiple packets are in turn grouped together to form what are referred to as a "sub-frame," and multiple sub-frames are in turn grouped together to form a "media" frame. To control the formation of frames and transfer of sub-frames, each sub-frame begins with a specially predefined packet referred to as a Sub-frame Header Packet (SHP).

[00199] The host device selects the data rate to be used for a given transfer. This rate can be changed dynamically by the host device based on both the maximum transfer capability of the host, or the data being retrieved from a source by the host, and the maximum capability of the client, or other device the data is being transferred to.

[00200] A recipient client device designed for, or capable of, working with the MDDI or inventive signal protocol is able to be queried by the host to determine the maximum, or current maximum, data transfer rate it can use, or a default slower minimum rate may be used, as well as useable data types and features supported. This information could be transferred using a Client Capability Packet (CCP), as discussed further below. The

client display device is capable of transferring data or communicating with other devices using the interface at a pre-selected minimum data rate or within a minimum data rate range, and the host will perform a query using a data rate within this range to determine the full capabilities of the client devices.

[00201] Other status information defining the nature of the bitmap and video frame-rate capabilities of the client can be transferred in a status packet to the host so that the host can configure the interface to be as efficient or optimal as practical, or desired within any system constraints.

[00202] The host sends filler packets when there are no (more) data packets to be transferred in the present sub-frame, or when the host cannot transfer at a rate sufficient to keep pace with the data transmission rate chosen for the forward link. Since each sub-frame begins with a sub-frame header packet, the end of the previous sub-frame contains a packet (most likely a filler packet) that exactly fills the previous sub-frame. In the case of a lack of room for data bearing packets per se, a filler packet will most likely be the last packet in a sub-frame, or at the end of a next previous sub-frame and before a sub-frame header packet. It is the task of the control operations in a host device to ensure that there is sufficient space remaining in a sub-frame for each packet to be transmitted within that sub-frame. At the same time, once a host device initiates the sending of a data packet, the host must be able to successfully complete a packet of that size within a frame without incurring a data under-run condition.

[00203] In one aspect of embodiments, sub-frame transmission has two modes. One mode is a periodic sub-frame mode, or periodic timing epochs, used to transmit live video and audio streams. In this mode, the Sub-frame length is defined as being non-zero. The second mode is an asynchronous or non-periodic mode in which frames are used to provide bitmap data to a client when new information is available. This mode is defined by setting the sub-frame length to zero in the Sub-frame Header Packet. When using the periodic mode, sub-frame packet reception may commence when the client has synchronized to the forward link frame structure. This corresponds to the "in sync" states defined according to the state diagram discussed below with respect to FIG. 49 or FIG. 63. In the asynchronous non-periodic sub-frame mode, reception commences after the first Sub-frame Header packet is received.

B. Overall Packet Structure

[00204] The format or structure of packets used to formulate the communication or signal protocol, or method or means for transferring data, implemented by the embodiments are presented below, keeping in mind that the interface is extensible and additional packet structures can be added as desired. The packets are labeled as, or divided into, different "packet types" in terms of their function in the interface, that is, commands, information, value, or data they transfer or are associated with. Therefore, each packet type denotes a pre-defined packet structure for a given packet which is used in manipulating the packets and data being transferred. As will be readily apparent, the packets may have pre-selected lengths or have variable or dynamically changeable lengths depending on their respective functions. The packets could also bear differing names, although the same function is still realized, as can occur when protocols are changed during acceptance into a standard. The bytes or byte values used in the various packets are configured as multi-bit (8- or 16-bit) unsigned integers. A summary of the packets being employed along with their "type" designations, listed in type order, is shown in Tables VI-1 through VI-4.

[00205] Each table represents a general "type" of packet within the overall packet structure for ease in illustration and understanding. There is no limitation or other impact implied or being expressed for the invention by these groupings, and the packets can be organized in many other fashions as desired. The direction in which transfer of a packet is considered valid is also noted.

Table VI -1
Link Control Packets

Packet Name	Packet Type	Valid in Forward	Valid in Reverse
Sub-frame Header Packet	15359	x	
Filler Packet	0	x	x
Reverse Link Encapsulation Packet	65	x	
Link Shutdown Packet	69	x	
Display Power State Packet	75	x	
Forward Audio Channel Enable Packet	78	x	
Perform Type Handoff Packet	77	x	
Round Trip Delay Measurement Packet	82	x	
Forward Link Skew Calibration Packet	83	x	

Table VI-2
Basic Media Stream Packets

Packet Name	Packet Type	Valid in Forward	Valid in Reverse
Video Stream Packet	16	x	x
Audio Stream Packet	32	x	x
Reserved Stream Packets	1 – 15, 18 – 31, 33 – 55	x	x
User-Defined Stream Packets	56 – 63	x	x
Color Map Packet	64	x	x
Reverse Audio Sample Rate Packet	79	x	
Transparent Color Enable Packet	81	x	

Table VI-3
Client Status and Control Packets

Packet Name	Packet Type	Valid in Forward	Valid in Reverse
Client Capability Packet	66		x
Keyboard Data Packet	67	x	x
Pointing Device Data Packet	68	x	x
Client Request and Status Packet	70		x
Digital Content Protection Overhead Packet	80	x	x
Request VCP Feature Packet	128	x	
VCP Feature Reply Packet	129		x
Set VCP Feature Packet	130	x	
Request Valid Parameter Packet	131	x	
Valid Parameter Reply Packet	132		x
Request Specific Status Packet	138	x	
Valid Status Reply List Packet	139		x
Personal Display Capability Packet	141		x
Client Error Report Packet	142		x
Scaled Video Stream Capability Packet	143		x
Client Identification Packet	144		x
Alternate Display Capability Packet	145		x
Register Access Packet	146	x	x

Table VI -4
Advanced Graphic and Display Packets

Packet Name	Packet Type	Valid in Forward	Valid in Reverse
Bitmap Block Transfer Packet	71	x	
Bitmap Area Fill Packet	72	x	
Bitmap Pattern Fill Packet	73	x	
Read Frame Buffer Packet	74	x	
Scaled Video Stream Capability Packet	143		x
Scaled Video Stream Setup Packet	136	x	
Scaled Video Stream Acknowledgement Packet	137		x
Scaled Video Stream Packet	18	x	

[00206] Something that is clear from other discussions within this text is that the Sub-frame Header, Filler, Reverse Encapsulation, Link Shutdown, Client Capability, and Client Request and Status Packets are each considered very important to, or even required in many embodiments of communication interfaces for External Mode operation. However Reverse Encapsulation, Link Shutdown, Client Capability, and Client Request and Status Packets, can be or are more likely to be considered optional for Internal Mode operation. This creates yet another type of MDDI protocol which allows communication of data at very high speeds with a reduced set of communications packets, and corresponding simplification of control and timing.

[00207] Packets have a common basic structure or overall set of minimum fields comprising a Packet Length field, a Packet Type field, Data Bytes field(s), and a CRC field, which is illustrated in FIG. 8. As shown in FIG. 8, the Packet Length field contains information, in the form of a multi-bit or -byte value, that specifies the total number of bits in the packet, or its length between the packet length field and the CRC field. In one embodiment, the packet length field contains a 16-bit or 2-byte wide, unsigned integer, that specifies the packet length. The Packet Type field is another multi-bit field which designates the type of information that is contained within the packet. In an exemplary embodiment, this is an 16-bit or 2-byte wide value, in the form of an 16-bit unsigned integer, and specifies such data types as display capabilities, handoff, video or audio streams, status, and so forth.

[00208] A third field is the Data Bytes field, which contains the bits or data being transferred or sent between the host and client devices as part of that packet. The format

of the data is defined specifically for each packet type according to the specific type of data being transferred, and may be separated into a series of additional fields, each with its own format requirements. That is, each packet type will have a defined format for this portion or field. The last field is the CRC field which contains the results of a 16-bit cyclic redundancy check calculated over the Data Bytes, Packet Type, and Packet Length fields, which is used to confirm the integrity of the information in the packet. In other words, calculated over the entire packet except for the CRC field itself. The client generally keeps a total count of the CRC errors detected, and reports this count back to the host in the Client Request and Status Packet (see further below).

[00209] Generally, these field widths and organization are designed to keep 2-byte fields aligned on an even byte boundary, and 4-byte fields aligned on 4-byte boundaries. This allows packet structures to be easily built in a main memory space of, or associated with, a host and a client without violating the data-type alignment rules encountered for most or typically used processors or control circuits.

[00210] During transfer of the packets, fields are transmitted starting with the Least Significant Bit (LSB) first and ending with the Most Significant Bit (MSB) transmitted last. Parameters that are more than one byte in length are transmitted using the least significant byte first, which results in the same bit transmission pattern being used for a parameter greater than 8 bits in length, as is used for a shorter parameter where the LSB is transmitted first. The data fields of each packet are generally transmitted in the order that they are defined in the subsequent sections below, with the first field listed being transmitted first, and the last field described being transmitted last. The data on the MDDI_Data0 signal path is aligned with bit '0' of bytes transmitted on the interface in any of the modes, Type 1, Type 2, Type 3, or Type-4.

[00211] When manipulating data for displays, the data for arrays of pixels are transmitted by rows first, then columns, as is traditionally done in the electronics arts. In other words, all pixels that appear in the same row in a bit map are transmitted in order with the left-most pixel transmitted first and the right-most pixel transmitted last. After the right-most pixel of a row is transmitted then the next pixel in the sequence is the left-most pixel of the following row. Rows of pixels are generally transmitted in order from top to bottom for most displays, although other configurations can be accommodated as needed. Furthermore, in handling bitmaps, the conventional approach, which is followed here, is to define a reference point by labeling the upper-left corner of a bitmap

as location or offset "0,0." The X and Y coordinates used to define or determine a position in the bitmap increase in value as one approaches the right and bottom of the bitmap, respectively. The first row and first column (upper left corner of an image) start with an index value of zero. The magnitude of the X coordinate increases toward the right side of the image, and the magnitude of the Y coordinate increases toward the bottom of the image as viewed by the user of the display.

[00212] A display window is the visible portion of a bitmap, the portion of the pixels in the bitmap that can be seen by the user on the physical display medium. It is often the case that the display window and the bitmap are the same size. The upper-left corner of a display window always displays bitmap pixel location '0,0'. The width of the display window corresponds to the X axis of the bitmap, and the display window width for this embodiment is less than or equal to the width of the corresponding bitmap. The height of the window corresponds to the Y axis of the bitmap, and the display window height for this embodiment is less than or equal to the height of the corresponding bitmap. The display window itself is not addressable in the protocol because it is only defined as the visible portion of a bitmap.

[00213] The relationship between a bitmaps and display windows is well known in the computer, electronic art, Internet communication, and other electronics related arts. Therefore, no further discussion or illustration of these principles is provided here.

C. Packet Definitions

1. Sub-Frame Header Packet

[00214] The Sub-Frame Header packet is the first packet of every sub-frame, and has a basic structure as illustrated in FIG. 9. The Sub-Frame Header Packet is used for host-client synchronization, every host should be able to generate this packet, while every client should be able to receive and interpret this packet. As can be seen in one embodiment in FIG. 9, this type of packet is structured to have Packet Length, Packet Type, Unique Word, Reserved 1, Sub-Frame Length, Protocol Version, Sub-Frame Count, and Media-frame Count fields, generally in that order. In one embodiment, this type of packet is generally identified as a Type 15359 (0x3bff hexadecimal) packet and uses a pre-selected fixed length of 20 bytes, not including the packet length field.

[00215] The Packet Type field and the Unique Word field each use a 2 byte value (16-bit unsigned integer). The 4-byte combination of these two fields together forms a 32-bit

unique word with good autocorrelation. In one embodiment, the actual unique word is 0x005a3bff, where the lower 16 bits are transmitted first as the Packet Type, and the most significant 16 bits are transmitted afterward.

[00216] The Reserved 1 field contains 2 bytes that are reserved space for future use, and is generally configured at this point with all bits set to zero. One purpose of this field is to cause subsequent 2-byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address. The least significant byte is reserved to indicate whether or not a host is capable of addressing multiple client devices. A value of zero for this byte is reserved to indicate that the host is capable of operating only with a single client device.

[00217] The Sub-frame Length field contains 4 bytes of information, or values, that specifies the number of bytes per sub-frame. In one embodiment, the length of this field may be set equal to zero to indicate that only one sub-frame will be transmitted by the host before the link is shut down into an idle state. The value in this field can be dynamically changed "on-the-fly" when transitioning from one sub-frame to the next. This capability is useful in order to make minor timing adjustments in the sync pulses for accommodating isochronous data streams. If the CRC of the Sub-frame Header packet is not valid then the link controller should use the Sub-frame Length of the previous known-good Sub-frame Header packet to estimate the length of the current sub-frame.

[00218] The Protocol Version field contains 2 bytes that specify the protocol version used by the host. The Protocol Version field may be set to '0' to specify the first or current version of the protocol as being used. This value will change over time as new versions are created, and is already being upgraded to a value of '1' for some version fields. Version values will probably or generally follow a current version number for an approved standards document which covers interfaces such as MDDI, as would be known.

[00219] The Sub-frame Count field contains 2 bytes that specify a sequence number that indicates the number of sub-frames that have been transmitted since the beginning of the media-frame. The first sub-frame of the media-frame has a Sub-frame Count of zero. The last sub-frame of the media-frame has a value of n-1, where n is the number of sub-frames per media-frame. The value of the Sub-frame Count field is equal to the Sub-frame Count sent in the previous Sub-Frame packet plus 1, except for a first sub-frame

of a media-frame which will have a count of zero. Note that if the Sub-frame Length is set equal to zero (indicating a non-periodic sub-frame) then the Sub-frame count is also set equal to zero.

- [00220] The Media-frame Count field contains 4 bytes (32-bit unsigned integer) that specify a sequence number that indicates the number of media-frames that have been transmitted since the beginning of the present media item or data being transferred. The first media-frame of a media item has a Media-frame Count of zero. The Media-frame Count increments just prior to the first sub-frame of each media-frame and wraps back to zero after the maximum Media-frame Count (for example, media-frame number $2^{31}-1 = 4,294,967,295$) is used. The Media-frame Count value may be reset generally at any time by the Host to suit the needs of an end application.

2. *Filler Packet*

- [00221] A filler packet is a packet that is transferred to, or from, a client device when no other information is available to be sent on either the forward or reverse link. It is recommended that filler packets have a minimum length in order to allow maximum flexibility in sending other packets when required. At the very end of a sub-frame or a reverse link encapsulation packet (see below), a link controller sets the size of the filler packet to fill the remaining space to maintain packet integrity. The Filler Packet is useful to maintain timing on the link when the host or client have no information to send or exchange. Every host and client needs to be able to send and receive this packet to make effective use of the interface.

- [00222] An exemplary embodiment of the format and contents of a Filler Packet are shown in FIG. 10. As shown in FIG. 10, this type of packet is structured to have Packet Length, Packet Type, Filler Bytes, and CRC fields. In one embodiment, this type of packet is generally identified as a Type 0, which is indicated in the 2-byte Type field. The bits or bytes in the Filler Bytes field comprise a variable number of all zero bit values to allow the filler packet to be the desired length. The smallest filler packet contains no bytes in this field. That is, the packet consists of only the packet length, packet type, and CRC, and in one embodiment uses a pre-selected fixed length of 6 bytes or a Packet Length value of 4. The CRC value is determined for all bytes in the packet including the Packet Length, which may be excluded in some other packet types.

3. *Video Stream Packet*

[00223] Video Stream Packets carry video data to update typically rectangular regions of a display device. The size of this region may be as small as a single pixel or as large as the entire display. There may be an almost unlimited number of streams displayed simultaneously, limited by system resources, because all context required to display a stream is contained within the Video Stream Packet. The format of one embodiment of a Video Stream Packet (Video Data Format Descriptor) is shown in FIG. 11. As seen in FIG. 11, in one embodiment, this type of packet is structured to have Packet Length (2 bytes), Packet Type, bClient ID, Video Data Descriptor, Pixel Display Attributes, X Left Edge, Y Top Edge, X Right Edge, Y Bottom Edge, X and Y Start, Pixel Count, Parameter CRC, Pixel Data, and Pixel Data CRC fields. This type of packet is generally identified as a Type 16, which is indicated in the 2-byte Type field. In one embodiment, a client indicates an ability to receive a Video Stream Packet using Red-Green-Blue (RGB), Monochrome, and Y Cr Cb Capability fields of the Client Capability Packet.

[00224] In one embodiment, the bClient ID field contains 2 bytes of information that are reserved for a Client ID. Since this is a newly developed communications protocol actual client IDs are not yet known or sufficiently communicable. Therefore, the bits in this field are generally set equal to zero until such ID values are known, at which time the ID values can be inserted or used, as would be apparent to those skilled in the art. The same process will generally be true for the client ID fields discussed below.

[00225] The format and contents employed for realizing the operation of an exemplary Video Data Descriptor field, as mentioned above, are shown in FIGS. 12A-12E. In FIGS. 12A-12E, the Video Data Format Descriptor field contains 2 bytes in the form of a 16-bit unsigned integer that specifies the format of each pixel in the Pixel Data in the present stream in the present packet. It is possible that different Video Stream packets may use different pixel data formats, that is, use a different value in the Video Data Format Descriptor, and similarly, a stream (region of the display) may change its data format on-the-fly. The pixel data format should comply with at least one of the valid formats for the client as defined in the Client Capability Packet. The Video Data Format Descriptor defines the pixel format for the present packet only which does not imply that a constant format will continue to be used for the lifetime of a particular video stream.

[00226] FIGs. 12A through 12D illustrate how the Video Data Format Descriptor is coded. As used in these figures, and in this embodiment, when bits [15:13] are equal to '000', as shown in FIG. 12A, then the video data consists of an array of monochrome pixels where the number of bits per pixel is defined by bits 3 through 0 of the Video Data Format Descriptor word. Bits 11 through 4 are generally reserved for future use or applications and are set to zero in this situation. When bits [15:13] are instead equal to the values '001', as shown in FIG. 12B, then the video data consists of an array of color pixels that each specify a color through a color map (palette). In this situation, bits 5 through 0 of the Video Data Format Descriptor word define the number of bits per pixel, and bits 11 through 6 are generally reserved for future use or applications and set equal to zero. When bits [15:13] are instead equal to the values '010', as shown in FIG. 12C, then the video data consists of an array of color pixels where the number of bits per pixel of red is defined by bits 11 through 8, the number of bits per pixel of green is defined by bits 7 through 4, and the number of bits per pixel of blue is defined by bits 3 through 0. In this situation, the total number of bits in each pixel is the sum of the number of bits used for red, green, and blue.

[00227] However, when bits [15:13] are instead equal to the values or string '011', as shown in FIG. 12D, then the video data consists of an array of video data in 4:2:2 YCbCr format with luminance and chrominance information, where the number of bits per pixel of luminance (Y) is defined by bits 11 through 8, the number of bits of the Cb component is defined by bits 7 through 4, and the number of bits of the Cr component is defined by bits 3 through 0. The total number of bits in each pixel is the sum of the number of bits used for red, green, and blue. The Cb and Cr components are sent at half the rate as Y. In addition, the video samples in the Pixel Data portion of this packet are organized as follows: Cbn, Yn, Cm, Yn+1, Cbn+2, Yn+2, Cm+2, Yn+3, ... where Cbn and Cm are associated with Yn and Yn+1, and Cbn+2 and Cm+2 are associated with Yn+2 and Yn+3, and so on.

[00228] Yn, Yn+1, Yn+2 and Yn+3 are luminance values of four consecutive pixels in a single row from left to right. If there are an odd number of pixels in a row ($X \text{ Right Edge} - X \text{ Left Edge} + 1$) in the window referenced by the Video Stream Packet then the Y value corresponding to the last pixel in each row will be followed by the Cb value of the first pixel of the next row, and a Cr value is not sent for the last pixel in the row. It is recommended that windows using Y Cb Cr format have a width that is an even

number of pixels. The Pixel Data in a packet should contain an even number of pixels. It may contain an odd or even number of pixels in the case where the last pixel of the Pixel Data corresponds to the last pixel of a row in the window specified in the Video Stream Packet header, i.e. when the X location of the last pixel in the Pixel Data is equal to X Right Edge.

[00229] When bits [15:13] are instead equal to the values '100' then the video data consists of an array of Bayer pixels where the number of bits per pixel is defined by bits 3 through 0 of the Video Data Format Descriptor word. The Pixel Group Pattern is defined by bits 5 and 4 as shown in FIG. 12E. The order of pixel data may be horizontal or vertical, and the pixels in rows or columns may be sent in forward or backward order and is defined by bits 8 through 6. Bits 11 through 9 should be set to zero. The group of four pixels in the pixel group in the Bayer format resembles what is often referred to as a single pixel in some display technologies. However, one pixel in the Bayer format is only one of the four colored pixels of the pixel group mosaic pattern.

[00230] For all five formats shown in the figures, Bit 12, which is designated as "P," specifies whether or not the Pixel Data samples are packed, or byte-aligned pixel data. A value of '0' in this field indicates that each pixel in the Pixel Data field is byte-aligned with an MDDI byte boundary. A value of '1' indicates that each pixel and each color within each pixel in the Pixel Data is packed up against the previous pixel or color within a pixel leaving no unused bits. The difference between Byte-Aligned and Packed Pixel data format is shown in more detail in FIG. 13, where one can clearly see that byte-aligned data may leave unused portions of the data sub-frame, as opposed to packed pixel format which does not.

4. Audio Stream Packet

[00231] The audio stream packets carry audio data to be played through the audio system of the client, or for a stand alone audio presentation device. Different audio data streams may be allocated for separate audio channels in a sound system, for example: left-front, right-front, center, left-rear, and right-rear, depending on the type of audio system being used. A full complement of audio channels is provided for headsets that contain enhanced spatial-acoustic signal processing. A client indicates an ability to receive an Audio Stream Packet using the Audio Channel Capability and Audio Sample

Rate fields of the Client Capability Packet. The format of Audio Stream Packets is illustrated in FIG. 14.

[00232] As shown in FIG. 14, this type of packet is structured in one embodiment to have Packet Length, Packet Type, bClient ID, Audio Channel ID, Reserved 1, Audio Sample Count, Bits Per Sample and Packing, Audio Sample Rate, Parameter CRC, Digital Audio Data, and Audio Data CRC fields. In one embodiment, this type of packet is generally identified as a Type 32 packet.

[00233] The bClient ID field contains 2 bytes of information that are reserved for a Client ID, as used previously. The Reserved 1 field contains 2 bytes that is reserved for future use, and is generally configured at this point with all bits set to zero.

[00234] The Bits Per Sample and Packing field contains 1 byte in the form of an 8-bit unsigned integer that specifies the packing format of audio data. The format generally employed is for Bits 4 through 0 to define the number of bits per PCM audio sample. Bit 5 then specifies whether or not the Digital Audio Data samples are packed. The difference between packed and byte-aligned audio samples, here using 10-bit samples, is illustrated in FIG. 15. A value of '0' indicates that each PCM audio sample in the Digital Audio Data field is byte-aligned with an MDDI byte boundary, and a value of '1' indicates that each successive PCM audio sample is packed up against the previous audio sample. This bit is generally effective only when the value defined in bits 4 through 0 (the number of bits per PCM audio sample) is not a multiple of eight. Bits 7 through 6 are reserved for future use and are generally set at a value of zero.

5. *Reserved Stream Packets*

[00235] In one embodiment, packet types 1 to 15, 18 to 31, and 33 through 55 are reserved for stream packets to be defined for use in future versions or variations of the packet protocols, as desired for various applications encountered. Again, this is part of making the MDDI more flexible and useful in the face of ever changing technology and system designs as compared to other techniques.

6. *User-Defined Stream Packets*

[00236] Eight data stream types, known as Types 56 through 63, are reserved for use in proprietary applications that may be defined by equipment manufacturers for use with a MDDI link. These are known as User-defined Stream Packets. Such packets may be

used for any purpose, but the host and client should only employ such packets in situations where the result of such use is very well understood or known. The specific definition of the stream parameters and data for these packet types is left to the specific equipment manufacturers or interface designers implementing such packet types or seeking their use. Some exemplary uses of the User-defined Stream Packets are to convey test parameters and test results, factory calibration data, and proprietary special use data. The format of the user-defined stream packets as used in one embodiment is illustrated in FIG. 16. As shown in FIG. 16, this type of packet is structured to have Packet Length (2 bytes), Packet Type, bClient ID number, Stream Parameters, Parameter CRC, Stream Data, and Stream Data CRC fields.

7. *Color Map Packets*

[00237] The color map packets specify the contents of a color map look-up table used to present colors for a client. Some applications may require a color map that is larger than the amount of data that can be transmitted in a single packet. In these cases, multiple Color Map packets may be transferred, each with a different subset of the color map by using the offset and length fields described below. The format of the Color Map Packet in one embodiment is illustrated in FIG. 17. As shown in FIG. 17, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Color Map Item Count, Color Map Offset, Parameter CRC, Color Map Data, and Data CRC fields. In one embodiment, this type of packet is generally identified as a Type 64 packet (Video Data Format and Color Map Packet) as specified in the Packet Type Field (2 bytes). A client indicates an ability to receive Color Map Packets using the Color Map Size and Color Map Width fields of the Client Capability Packet.

8. *Reverse Link Encapsulation Packets*

[00238] In an exemplary embodiment, data is transferred in the reverse direction using a Reverse Link Encapsulation Packet. A forward link packet is sent and the MDDI link operation (transfer direction) is changed or turned around in the middle of this packet so that packets can be sent in the reverse direction. The format of the Reverse Link Encapsulation packet in one embodiment is illustrated in FIG. 18. As shown in FIG. 187, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Reverse Link Flags, Reverse Rate Divisor, Turn-Around 1 Length, Turn-Around 2

Length, Parameter CRC, All Zero 1, Turn-Around 1, Reverse Data Packets, Turn-Around 2, and All Zero 2 fields. In one embodiment, this type of packet is generally identified as a Type 65 packet. For External Mode every host must be able to generate this packet and receive data, and every client must be able to receive and send data to the host in order to efficiently make use of the desired protocol and resulting speed. Implementation of this packet is optional for Internal Mode, but the Reverse Link Encapsulation Packet is used for the host to receive data from the client.

[00239] The MDDI link controller behaves in a special manner while sending a Reverse Link Encapsulation Packet. The MDDI has a strobe signal that is generally always driven by the host as controller of the link. The host behaves as if it were transmitting a zero for each bit of the Turn-Around and Reverse Data Packets portions of the Reverse Link Encapsulation packet. The host toggles a MDDI_Strobe signal at each bit boundary during the two turn-around times and during the time allocated for reverse data packets. That is, the host toggles MDDI_Stb from the beginning of the All Zero 1 field to the end of the All Zero 2 field. (This is the same behavior as if it were transmitting all-zero data.)

[00240] The host disables its MDDI data signal line drivers and generally assures they have been completely disabled prior to the last bit of the Turn-Around 1 field, and then re-enables its line drivers during the Turn-Around 2 field, and generally assure that the drivers have been completely re-enabled prior to the last bit of the Turn-Around 2 field. The client reads the Turn-Around Length parameter and drives the data signals toward the host immediately after the last bit in the Turn-Around 1 field. That is, the client clocks new data into the link on certain rising edges of the MDDI strobe as specified in the packet contents description below, and elsewhere. The client uses the Packet Length and Turn-Around Length parameters to know the length of time it has available to send packets to the host. The client may send filler packets or drive the data lines to a zero state when it has no data to send to the host. If the data lines are driven to zero, the host interprets this as a packet with a zero length (not a valid length) and the host does not accept any more packets from the client for the duration of the current Reverse Link Encapsulation Packet.

[00241] In one embodiment, the Reverse Link Request field of the Client Request and Status Packet may be used to inform the host of the number of bytes the client needs in the Reverse Link Encapsulation Packet to send data back to the host. The host attempts

to grant the request by allocating at least that number of bytes in the Reverse Link Encapsulation Packet. The host may send more than one Reverse Link Encapsulation Packet in a sub-frame. The client may send a Client Request and Status Packet at almost any time, and the host will interpret the Reverse Link Request parameter as the total number of bytes requested in one sub-frame.

9. *Client Capability Packets*

[00242] A host needs to know the capability of the client (display) it is communicating with in order to configure the host-to-client link in an generally optimum or desired manner. It is recommended that a display send a Client Capability Packet to the host after forward link synchronization is acquired. The transmission of such a packet is considered required when requested by the host using the Reverse Link Flags in the Reverse Link Encapsulation Packet. The Client Capability Packet is used to inform the host of the capabilities of a client. For External Mode every host should be able to receive this packet, and every client should be able to send this packet to fully utilize this interface and protocol. Implementation of this packet is optional for Internal Mode, since the capabilities of the client, such as a display, keyboard or other input/output device, in this situation should already be well defined and known to the host at the time of manufacture or assembly into a single component or unit of some type.

[00243] The format of the Client Capability packet in one embodiment is illustrated in FIG. 19. As shown in FIG. 19, for this embodiment, this type of packet is structured to have Packet Length, Packet Type, cClientID, Protocol Version, Min Protocol Version, Data Rate Capability, Interface Type Capability, Number of Alt Displays, Reserved 1, Bitmap Width, Bitmap Height, Display Window Width, Display Window Height, Color Map Size, Color Map RGB Width, RGB Capability, Monochrome Capability, Reserved 2, Y Cr Cb Capability, Bayer Capability, Reserved 3, Client Feature Capability, Max Video Frame Rate, Min Video Frame Rate, Min Sub-frame rate, Audio Buffer Depth, Audio Channel Capability, Audio Sample Rate Capability, Audio Sample Resolution, Mic Sample Resolution, Mic Sample Rate Capability, Keyboard Data Format, Pointing Device Data Format, Content Protection Type, Mfr. Name, Product Code, Reserved 4, Serial Number, Week of Mfr., Year of Mfr., and CRC fields. In an exemplary embodiment, this type of packet is generally identified as a Type 66 packet.

10. Keyboard Data Packets

[00244] A keyboard data packet is used to send keyboard data from the client device to the host. A wireless (or wired) keyboard may be used in conjunction with various displays or audio devices, including, but not limited to, a head mounted video display/audio presentation device. The Keyboard Data Packet relays keyboard data received from one of several known keyboard-like devices to the host. This packet can also be used on the forward link to send data to the keyboard. A client indicates an ability to send and receive Keyboard Data Packets using the Keyboard Data Field in the Client Capability Packet.

[00245] The format of a Keyboard Data Packet is shown in FIG. 20, and contains a variable number of bytes of information from or for a keyboard. As shown in FIG. 20, this type of packet is structured to have Packet Length, Packet Type, bClient ID, Keyboard Data Format, Keyboard Data, and CRC fields. Here, this type of packet is generally identified as a Type 67 packet.

[00246] The bClient ID is a reserved field, as before, and the CRC is performed over all bytes of the packet. The Keyboard Data Format field contains a 2 bytes value that describes the keyboard data format. Bits 6 through 0 should be identical to the Keyboard Data Format field in the Client Capability Packet. This value is not to equal 127. Bits 15 through 7 are reserved for future use and are, therefore, currently set to zero.

11. Pointing Device Data Packets

[00247] A pointing device data packet is used as a method, structure, or means to send position information from a wireless mouse or other pointing device from the client to the host. Data can also be sent to the pointing device on the forward link using this packet. An exemplary format of a Pointing Device Data Packet is shown in FIG. 21, and contains a variable number of bytes of information from or for a pointing device. As shown in FIG. 21, this type of packet is structured to have Packet Length, Packet Type, bClient ID, Pointing Device Format, Pointing Device Data, and CRC fields. In an exemplary embodiment, this type of packet is generally identified as a Type 68 packet in the 1-byte type field.

12. *Link Shutdown Packets*

[00248] A Link Shutdown Packet is sent from the host to the client as a method or means to indicate that the MDDI data and strobe will be shut down and go into a low-power consumption "hibernation" state. This packet is useful to shut down the link and conserve power after static bitmaps are sent from a mobile communication device to the display, or when there is no further information to transfer from a host to a client for the time being. Normal operation is resumed when the host sends packets again. The first packet sent after hibernation is a sub-frame header packet. The format of a Client Status Packet for one embodiment is shown in FIG. 22. As shown in FIG. 22, this type of packet is structured to have Packet Length, Packet Type, CRC and All Zeros fields. In one embodiment, this type of packet is generally identified as a Type 69 packet in the 1-byte type field.

[00249] The packet length field uses 2 bytes to specify the total number of bytes in the packet not including the packet length field. In one embodiment, the Packet Length of this packet is dependent on the Interface Type or link mode in effect at the time when the Link Shutdown Packet is sent. Therefore, the typical packet length becomes 20 bytes for Type 1 mode (22 bytes total in the packet), 36 bytes for a Type 2 mode (38 bytes total in the packet), 68 bytes for a Type 3 mode link (70 bytes total in the packet), and 132 bytes for a Type 4 mode (with 134 bytes total in the packet).

[00250] The All Zeros field uses a variable number of bytes to ensure that MDDI_Data signals are at a logic-zero level for a sufficient time to allow the client to begin recovering clock using only MDDI_Stb prior to disabling a host's line drivers. The length of the All Zeros field is dependent on the Interface Type or link operating mode in effect at the time when the Link Shutdown Packet is sent. The length of the All Zeros field is intended to produce 64 pulses on MDDI_Stb for any Interface Type setting. Therefore, the All Zeros length for each interface type becomes 16 bytes for Type 1, 32 bytes for Type 2, 64 bytes for Type 3, and 128 bytes for Type 4.

[00251] The CRC field uses 2 bytes that contain a 16-bit CRC of bytes from the Packet Length to the Packet Type.

[00252] In the low-power hibernation state, the MDDI_Data0 driver is disabled into a high-impedance state starting after the 16th to 48th MDDI_Stb cycle or pulse after the last bit of the All Zeros field. For Type-2, Type-3, or Type-4 links the MDDI_Data1 through MDDI_DataPwr7 signals are also placed in a high-impedance state at the same

time that the MDDI_Data0 driver is disabled Either the host or client may cause the MDDI link to "wake up" from the hibernation state as described elsewhere, which is a key advance for and advantage of the present invention.

[00253] As described in the definition of the All Zeros field, MDDI_Stb toggles for 64 cycles following the MSB of the CRC field of the Link Shutdown Packet to facilitate an orderly shutdown in the client controller. One cycle is a low-to-high transition followed by a high-to-low transition, or a high-to-low transition followed by a low-to-high transition. After the All Zeros field is sent, the MDDI_Stb driver in the host is disabled.

13. *Client Request and Status Packets*

[00254] The host needs a small amount of information from the client so it can configure the host-to-client link in a generally optimum manner. It is recommended that the client send one Client Request and Status Packet to the host each sub-frame. The client should send this packet as the first packet in the Reverse Link Encapsulation Packet to ensure that it is delivered reliably to the host. The forwarding of this packet is also accomplished when requested by a host using the Reverse Link Flags in the Reverse Link Encapsulation Packet. The Client Request and Status Packet is used to report errors and status to the host. For external mode operation, every host should be able to receive this packet, and every client should be able to send this packet in order to properly or optimally employ the MDDI protocol. While it is also recommended that for internal operations, that is internal hosts and internal clients, there should be support for this packet, it is not required.

[00255] The format of a Client Request and Status Packet is shown in FIG. 23. As shown in FIG. 23, this type of packet is structured to have Packet Length, Packet Type, cClient ID, Reverse Link Request, Capability Change, Client Busy, CRC Error Count, and CRC fields. This type of packet is generally identified as a Type 70 packet in the 1-byte type field, and typically uses a pre-selected fixed length of 12 bytes.

[00256] The Reverse Link Request field may be used to inform the host of the number of bytes the client needs in the Reverse Link Encapsulation Packet to send data back to the host. The host should attempt to grant the request by allocating at least that number of bytes in the Reverse Link Encapsulation Packet. The host may send more than one Reverse Link Encapsulation Packet in a sub-frame in order to accommodate data. The client may send a Client Request and Status Packet at any time and the host will

interpret the Reverse Link Request parameter as the total number of bytes requested in one sub-frame. Additional details and specific examples of how reverse link data is sent back to the host are shown below.

14. Bit Block Transfer Packets

[00257] The Bit Block Transfer Packet provides a means, structure, or method to scroll regions of the display in any direction, generally by copying a block of pixels from one rectangular region to another. Clients that have this capability will report the capability in bit 0 of the Display Feature Capability Indicators field of the Client Capability Packet. The format for one embodiment of a Bit Block Transfer Packet is shown in FIG. 24. As shown in FIG. 24, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Pixel Data Attributes, Raster Operation, Upper Left X Value, Upper Left Y Value, Window Width, Window Height, Window X Movement, Window Y Movement, and CRC fields. This type of packet is generally identified as a Type 71 packet, and in one embodiment uses a pre-selected fixed length of 15 bytes. The 2-byte hClient ID field contains information or values that are reserved for a Client ID, as discussed elsewhere. Since this field is generally reserved for future use, the current value is typically set to zero, by setting the bits to a logic-zero level, although it can be set to other values or used by one skilled in the art to transfer desired information.

[00258] In one embodiment, the 2-byte Pixel Data Attributes field has values that specify where the pixel data is going to be updated, with Bits 1 and 0 selecting the display where the pixel data is going to be updated. If a primary display in the client does not support stereo images then the client can affect the pixel data in the primary display for one of the bit combinations 01, 10, or 11. It is recommended that the value 11 be used to address the primary display in clients that do not support stereo display capability. When Bits [1:0] have the values 11, the pixel data is updated in the frame buffer of both the left and right eye, if Bits [1:0] have the values 10, the pixel data is updated in the frame buffer of the left eye only. When Bits [1:0] have the values 01, the pixel data is updated in the frame buffer of the right eye only. When Bits [1:0] have the values 00, the pixel data is updated in the frame buffer of the alternate display specified by bits 8 through 11 below.

[00259] Bits 7 and 6 act as Display Update Bits that specify the frame buffer where the pixel data is to be updated or written. The effects of the Frame Update Bits are

described in more detail later. When Bits [7:6] are '01', the Pixel data is written to an offline image buffer. When Bits [7:6] are '00', the Pixel data is written to an image buffer used to refresh the display. When Bits [7:6] are '11', the Pixel data is written to all image buffers. If Bits [7:6] are '10', this is treated as an invalid value. These bits are currently reserved for future use. In this situation, the entire command is ignored and no frame buffers are updated.

[00260] Bits 11 through 8 form a 4-bit unsigned integer that specifies an alternate display or alternative client location where the pixel data is to be updated. Bits 0 and 1 are set equal to 00 in order for a client to interpret bits 11 through 8 as an alternate display number. If bits 1 and 0 are not equal to 00 then bits 8 through 11 are generally set equal to a logic-zero value or level. Bits 2 through 5 and 12 through 15 are reserved for future use and are generally be set to logic-zero level or values.

[00261] In one embodiment, the 2 byte Raster Operation field specifies how to combine pixels in source and destination locations to form new pixel values to be written to a destination image location. Raster operations define how two different rectangular regions of equal size in a frame buffer are merged together. The destination image area is also one of the two images that are merged together. The second of the two images is called the source image. If the client does not support the Raster Operation field as specified in the Client Capability Packet then the host generally sends this packet with bits 3 through 0 equal to 3, and the client ignores bits 3 to 0.

[00262] In one embodiment, Bits 3 to 0 are used to specify an actual raster operation by using or setting them equal to one of the values shown in Table VII below to select the corresponding operation shown next to that value. That is, each specified Bits [3:0] value listed in the first column results in the operation specified in the second column, and further defined here for clarification in the third column.

Table VII

Bits [3:0] Value	Value stored in Destination Location	Definition
0	0	
1	source & destination	logic AND operation
2	source & ~destination	source AND (not dest)
3	source	
4	~source & destination	(not source) AND dest
5	destination	no operation
6	source ^ destination	logic XOR operation
7	source destination	logic OR operation
8	~(source destination)	not (source OR dest)
9	~(source ^ destination)	not (source XOR dest)
10	~(destination)	not (dest)
11	source ~destination	source OR (not dest)
12	~source	not source
13	~source destination	(not source) OR dest
14	~(source & destination)	not (source AND dest)
15	All ones	

[00263] Bits 5 through 4 are used to specify whether or not the destination pixels are written to the destination locations as they relate to the transparent color. The operation specified by bits 5 to 4 apply whether or not the raster operations is supported by the client device. If the client does not support raster operations then the resulting destination pixel value to be considered for the operation defined by bits 5 through 4 is equal to the source pixel value only.

[00264] When the value of Bits [5:4] is equal to 00, then transparent color is not used. A resulting destination pixel is written to the destination pixel location without considering the value of the transparent color defined by the Transparent Color Enable Packet. The value of Bits [5:4] being equal to 01 is currently reserved for future use and typically not used, although available for one skilled in the art to establish a related use for. When the value of Bits [5:4] is equal to 10, the resulting pixel is not written to the destination pixel location if the resulting destination pixel computed by the raster operation is equal to the transparent color. Otherwise it is written to the destination pixel location. When the value of Bits [5:4] is equal to 11 the resulting pixel is not written to the destination pixel location if the resulting destination pixel computed by

the raster operation is equal to the transparent color.. Otherwise the resulting pixel is not written to the destination pixel location.

[00265] Bits 15 to 6 are reserved for future use and are, therefore, generally set equal to a logic-zero value or level.

[00266] The remaining fields are used to specify the X and Y values of the coordinate of the upper left corner of the window to be moved, the width and height of the window to be moved, and the number of pixels that the window is to be moved horizontally, and vertically, respectively. Positive values for the latter two fields cause the window to be moved to the right, and down, and negative values cause movement to the left and up, respectively. The CRC field (here 2 bytes) contains a 16-bit CRC of all bytes in the packet including the Packet Length.

15. *Bitmap Area Fill Packets*

[00267] The Bitmap Area Fill Packet provides a means, structure, or method to easily initialize a region of the display to a single color. Displays that have this capability will report the capability in bit 1 of the Client Feature Capability Indicators field of the Client Capability Packet. One embodiment for the format of a Bitmap Area Fill Packet is shown in FIG. 25. As shown in FIG. 25, in this case this type of packet is structured to have Packet Length, Packet Type, hClient ID, Upper Left X Value, Upper Left Y Value, Window Width, Window Height, Data Format Descriptor, Pixel Area Fill Value, and CRC fields. This type of packet is generally identified as a Type 72 packet in the 2-byte type field, and uses a pre-selected fixed length of 20 bytes.

[00268] The 2-byte hClient ID field contains information or values that are reserved for a Client ID, as discussed elsewhere. Since this field is generally reserved for future use, the current value is typically set to zero, by setting the bits to a logic-zero level, although it can be set to other values or used by one skilled in the art to transfer desired information.

16. *Bitmap Pattern Fill Packets*

[00269] The Bitmap Pattern Fill Packet provides a means or structure to easily initialize a region of the display to a pre-selected pattern. Clients that have this capability will report the capability in bit 2 of the Client Feature Capability field of the Client Capability Packet. The upper left corner of the fill pattern is aligned with the upper left

corner of the window to be filled, unless the horizontal or vertical pattern offset is non-zero. If the window to be filled is wider or taller than the fill pattern, then the pattern may be repeated horizontally or vertically a number of times to fill the window. The right or bottom of the last repeated pattern is truncated as necessary. If the window is smaller than the fill pattern, then the right side or bottom of the fill pattern may be truncated to fit the window.

[00270] If a horizontal pattern offset is non-zero, then the pixels between the left side of the window and the left side plus the horizontal pattern offset are filled with the right-most pixels of the pattern. The horizontal pattern offset is to be less than the pattern width. Similarly, if a vertical pattern offset is non-zero, then the pixels between the top side of the window and the top of the side plus vertical pattern offset are filled with the lower-most pixels of the pattern. The vertical pattern offset is to be less than the pattern height.

[00271] One embodiment for the format of a Bitmap Pattern Fill Packet is shown in FIG. 26. As shown in FIG. 26, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Upper Left X Value, Upper Left Y Value, Window Width, Window Height, Pattern Width, Pattern Height, Horizontal Pattern Offset, Vertical Pattern Offset, Data Format Descriptor, Parameter CRC, Pattern Pixel Data, and Pixel Data CRC fields. In some embodiments, this type of packet is generally identified as a Type 73 packet in the 1-byte type field.

17. Communication Link Data Channel Packets

[00272] The Communication Link Data Channel Packet provides a structure, means, or method for a client with high-level computing capability, such as a PDA, to communicate with a wireless transceiver such as a cell phone or wireless data port device. In this situation, the MDDI link is acting as a convenient high-speed interface between the communication device and the computing device with the mobile display, where this packet transports data at a Data Link Layer of an operating system for the device. For example, this packet could be used if a web browser, email client, or an entire PDA were built into a mobile display. Displays that have this capability will report the capability in bit 3 of the Client Feature Capability field of the Client Capability Packet.

[00273] The format of an embodiment for a Communication Link Data Channel Packet is shown in FIG. 27. As shown in FIG. 27, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Parameter CRC, Communication Link Data, and Communication Data CRC fields. In one embodiment, this type of packet is generally identified as a Type 74 packet in the type field.

18. Display Power State Packets

[00274] The Display Power State Packet provides a structure, means, or method for placing specific client controlled or client related, connected, or controller hardware into a low power state when a client such as a display is not being used or in current active use, in order to minimize the system power consumption or drain on system resources. A packet of this type is the most useful for applications of the interface or interface structure and protocol to external mode configurations or operations. In such applications, it is more likely that the external device is operating on limited power resources such as batteries, or has other power constraints and concerns, for example overheating in limited spaces, and so forth, such that a minimal operating condition is desired for periods or inactivity or non-use. In one embodiment, a client indicates an ability to respond to Display Power State Packets using bit 9 of the Client Feature Capability Indicators field of the Client Capability Packet.

[00275] The format of one embodiment for a Display Power State Packet is shown in FIG. 28. As shown in FIG. 28, in one embodiment, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Power State, and CRC fields. This type of packet is generally identified as a Type 75 packet in the 2-byte type field. The 2-byte hClient ID field contains information or values that are reserved for a Client ID, as used previously. Since this field is generally reserved for future use, the current value is set to zero, by setting the bits to '0', although it can be used by one skilled in the art to transfer desired information.

[00276] The Power State field, here 2 bytes, specifies the information used to place a specific device, piece of hardware, or equipment associated with the client such as a display into the specified power state. When used for displays, Bit 0 of this field specifies whether or not the packet applies to the main display or to an alternate display. If bit 0 is equal to 1 then the packet applies to the main display. If bit 0 is equal to 0

then the packet applies to the alternate display specified by bits 11 through 8. Bit 1 is reserved for future use and is generally set to zero.

[00277] Bits 3 through 2 of the Power State field specify the power state of the display selected by bits 11 through 8 and bit 0. When Bits[3:2] have a value of '00', the selected display is not illuminated and should be consuming a minimum amount of power, and the contents of the frame buffer are not guaranteed to be retained during this state. When Bits[3:2] have a value of '01', the selected display is not illuminated and is consuming a relative minimum amount of power and the contents of the frame buffer are guaranteed to be retained during this state. The display may consume more power in this state than in state 00. The client can indicate an ability to support state 01 using bit 10 of the Client Feature Capability Indicators field of the Client Capability Packet. When Bits[3:2] of the Power State field have a value of '10', the selected display is illuminated and is displaying an image from its associated frame buffer. The value of '11' for Bits[3:2] is a reserved value or state for future use and is not used.

[00278] Those skilled in the art will recognize that while most useful for display applications, use of this packet is not limited by this invention to only displays and there may be other applications, configurations, or situations in which power control may be needed or desired in relation to other hardware elements with which the MDDI is being used, or for which a client is controlling or communicating. In these situations, the Bits disclosed above may have similar functions but could be activating main and secondary ones of such elements, or setting power levels and so forth, as would be understood.

[00279] In one embodiment, Bits 11 through 8 of the Power State field form a 4-bit unsigned integer that specifies the alternate display to which the power state is applied. Bit 0 is set to a logic-zero value in order for the client to interpret bits 11 through 8 as an alternate display number. If bit 0 is equal to 1 then bits 11 through 8 are zero.

[00280] Bits 7 through 4 and Bits 15 through 12 are reserved for future use, and are generally set to logic-zero level or values for current applications or designs.

[00281] The 2 byte CRC field specifies or contains the CRC of all bytes in the packet including the Packet Length.

[00282] A summary of which display power states are generally supported by the interface structure or protocol is presented in Table VIII below. As can be seen, various combinations of Client Feature Capability Bits 10 and 9 are used to establish, setup, or trigger various ones of the desired power states. A mark present in given row and

column position indicates that the display power state specified at the top of that column is supported for the stated combination of Client Feature Capability Indicator bits.

TABLE VIII

Client Feature Capability Indicator Bits 9 and 10 values	Power State = 00	Power State = 01	Power State = 10	Power State = 11
Bit 9 = 0 and Bit 10 = 0	X		X	X
Bit 9 = 1 and Bit 10 = 0			X	X
Bit 9 = 0 and Bit 10 = 1	X	X	X	X
Bit 9 = 1 and Bit 10 = 1			X	X

19. Perform Type Handoff Packets

[00283] The Perform Type Handoff Packet is a means, structure, or method for the host to use in order to command a client to handoff to the mode specified in this packet. This is to be one of the interface type settings supported by the client as described in the Client Capability Packet. The host and client should switch to the specified forward and reverse link interface type right after this packet is sent. The format of one embodiment for a Perform Type Handoff Packet is shown in FIG. 29. Hosts and clients that support an interface type other than Type 1 should provide support for this packet. It is typically recommended that a host reads the Client Request and Status Packet immediately before it sends the Perform Type Handoff Packet to confirm that the client is in sync with the host.

[00284] As shown in FIG. 29, in one embodiment, this type of packet is structured to have Packet Length, Packet Type, Interface Type, Reserve 1, Delay Filler, and CRC fields. This type of packet is generally identified as a Type 77 packet in the 2-byte type field, and uses a pre-selected fixed length of 6 bytes, outside of the Packet Length and Delay Filler fields.

[00285] In one embodiment, the Interface Type field uses a 1 byte value to confirm a new interface type to be used or employed for the link. The value in this field specifies or represents the interface type in the following manner. Bits 2 through 0 define the interface Type to be used on the forward link with a value of 1 signifying or specifying a handoff to a Type 1 mode; a value of 2 a handoff to Type 2 mode, a value of 3 a handoff to Type 3 mode, and a value of 4 a handoff to Type 4 mode. Bits 5 through 3

define the interface Type to be used on the reverse link with a value of 1 signifying or specifying a handoff to a Type 1 mode, value of 2 a handoff to Type 2 mode, a value of 3 a handoff to Type 3 mode, and a value of 4 a handoff to Type 4 mode. Bits 0, 6, and 7 are currently reserved for future use, and as such are typically, but not necessarily, set to a logic-zero level.

[00286] The Delay Filler field has been created as a means, structure, or method for allowing sufficient time on the part of the system for the client to prepare or be configured to switch over to use or set up for using a new interface type setting at the beginning of the packet that immediately follows the Perform Interface Type Handoff Packet. This field contains a group of bytes or 8-bit values that are all set at or equal to a logic-zero level or value. The number of bytes used in this field is selected such that it results in this field being a length equivalent to 64 MDDI_Stb cycles. The length of the Delay Filler field is based on the interface type setting of the forward link which will be 16 bytes for a Type 1 forward link interface type, 32 bytes for a Type 2 interface type, 64 bytes for a Type 3 interface type, and 128 bytes when specifying or using a Type 4 forward link interface type.

[00287] The Reserved 1 field (here 1 byte) is reserved for future use in imparting information. All bits in this field are generally set to a logic-zero level. The purpose of such fields is currently to cause all subsequent 2 byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address. The CRC field (here 2 bytes) contains a 16-bit CRC of all bytes in the packet including the Packet Length.

20. Forward Audio Channel Enable Packets

[00288] This packet provides a structure, method, or means that allows a host to enable or disable audio channels in a client. This capability is useful so that a client (a display for example) can power off audio amplifiers or similar circuit elements to save power when there is no audio to be output by the host. This is significantly more difficult to implement implicitly simply using the presence or absence of audio streams as an indicator. The default state when the client system is powered-up is that all audio channels are enabled. The format of one embodiment of a Forward Audio Channel Enable Packet is shown in FIG. 30. As shown in FIG. 30, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Audio Channel Enable

Mask, and CRC fields. This type of packet is generally identified as a Type 78 packet in the 1-byte type field, and uses a pre-selected fixed length of 4 bytes.

21. Reverse Audio Sample Rate Packets

[00289] This type of packet provides a structure, method, or means that allows a host to enable or disable audio channels in a client. This capability is useful so the client can power off audio amplifiers to save power when there is no audio to be output by the host. This is significantly more difficult to implement implicitly using the presence or absence of audio streams. The default state when a client system is powered-up or connected to the host is that all audio channels are enabled. An audio system connected to a host and a client should be ready or able to output audio signals in an intended or desired manner within about 100 msec. or less after the client receives a Forward Audio Channel Enable packet having a least one of the bits in the Audio Channel Enable Mask field having made a transition from a zero to a one state or value. The client indicates an ability to respond to a Forward Audio Channel Enable Packet using the value set for bit 15 of the Audio Channel Capability field of the Client Capability Packet.

[00290] This packet allows the host to enable or disable the reverse-link audio channel, and to set the audio data sample rate of this stream. The host selects a sample rate that is defined to be valid in the Client Capability Packet. If the host selects an invalid sample rate then the client will not send an audio stream to the host, and an appropriate error, error value, or error signal, may be sent to the host in the Client Error Report Packet. The host may disable the reverse-link audio stream by setting the sample rate to a value of 255. The default state assumed when the client system is initially powered-up or connected is with the reverse-link audio stream disabled. The format of one embodiment for a Reverse Audio Sample Rate Packet is shown in FIG. 31. As shown in FIG. 31, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Audio Sample Rate, Reserved 1, and CRC fields. This type of packet is generally identified as a Type 79 packet, and uses a pre-selected fixed length of 4 bytes.

22. Digital Content Protection Overhead Packets

[00291] This packet provides a structure, method, or means that allows a host and a client to exchange messages related to the digital content protection method being used. Presently two types of content protection are contemplated, Digital Transmission

Content Protection (DTCP), or High-bandwidth Digital Content Protection (HDCP) system, with room reserved for future alternative protection scheme designations. The method being used is specified by a Content Protection Type parameter in this packet. The format of an embodiment of a Digital Content Protection Overhead Packet is shown in FIG. 32. As shown in FIG. 32, this type of packet is structured to have Packet Length, Packet Type, bClient ID, Content Protection Type, Content Protection Overhead Messages, and CRC fields. This type of packet is generally identified as a Type 80 packet.

23. *Transparent Color Enable Packets*

[00292] The Transparent Color Enable Packet is a structure, method, or means that used to specify which color is transparent in a display and to enable or disable the use of a transparent color for displaying images. Displays that have this capability will report that capability in bit 4 of the Client Feature Capability field of the Client Capability Packet. When a pixel with the value for transparent color is written to the bitmap, the color does not change from the previous value. The format of a Transparent Color Enable Packet is shown in FIG. 33. As shown in FIG. 33, in one embodiment this type of packet is structured to have Packet Length, Packet Type, hClient ID, Transparent Color Enable, Reserved 1, Alpha-Cursor Identifier, Data Format Descriptor, Transparent Pixel Value, and CRC fields. This type of packet is generally identified as a Type 81 packet in the 1-byte type field, and uses a pre-selected fixed length of 10 bytes.

24. *Round Trip Delay Measurement Packets*

[00293] The Round Trip Delay Measurement Packet provides a structure, method, or means that is used to measure the propagation delay from the host to a client (display) plus the delay from the client (display) back to the host. This measurement inherently includes the delays that exist in the line drivers and receivers, and an interconnect subsystem. This measurement is used to set the turn around delay and reverse link rate divisor parameters in the Reverse Link Encapsulation Packet, described generally above. This packet is most useful when the MDDI link is running at the maximum speed intended for a particular application. The packet may be sent in Type 1 mode and at a lower data rate in order to increase the range of the round trip delay measurement.

The MDDI_Stb signal behaves as though all zero data is being sent during the following fields: both Guard Times, All Zero, and the Measurement Period. This causes MDDI_Stb to toggle at half the data rate so it can be used as periodic clock in the client during the Measurement Period.

[00294] In one embodiment, a client generally indicates an ability to support the Round Trip Delay Measurement Packet through use of bit 18 of the Client Feature Capability Indicators field of the Client Capability Packet. It is recommended that all clients support round trip delay measurement, but it is possible for the host to know the worst-case round trip delay based on a maximum cable delay, and on maximum driver and receiver delays. The host may also know the round-trip delay in advance for an MDDI link used in internal mode, since this is an aspect of known design elements (conductor lengths, circuitry type, and features, and so forth) of the device in which the interface is being used.

[00295] The format of a Round Trip Delay Measurement Packet is shown in FIG. 34. As shown in FIG. 34, in one embodiment this type of packet is structured to have Packet Length, Packet Type, hClient ID, Parameter CRC, Guard Time 1, Measurement Period, All Zero, and Guard Time 2 fields. This type of packet is generally identified as a Type 82 packet, and uses a pre-selected fixed length of 159 bits.

[00296] The timing of events that take place during the Round Trip Delay Measurement Packet is illustrated in FIG. 35. In FIG. 35, the host transmits the Round Trip Delay Measurement Packet, shown by the presence of the Parameter CRC and Strobe Alignment fields followed by the All Zero 1 and Guard Time 1 fields. A delay 3502 occurs before the packet reaches the client display device or processing circuitry. As the client receives the packet, it transmits the 0xff, 0xff, and 30 bytes of 0x00 pattern as precisely as practical at the beginning of the Measurement Period as determined by the client. The actual time the client begins to transmit this sequence is delayed from the beginning of the Measurement Period from the point of view of the host. The amount of this delay is substantially the time it takes for the packet to propagate through the line drivers and receivers and the interconnect subsystem (cables, conductors). A similar amount of delay 3504 is incurred for the pattern to propagate from the client back to the host.

[00297] In order to accurately determine the round trip delay time for signals traversing to and from the client, the host counts the number of forward link bit time periods

occurring after the start of the Measurement Period until the beginning of the 0xff, 0xff, and 30 bytes of 0x00 sequence is detected upon arrival. This information is used to determine the amount of time for a round trip signal to pass from the host to the client and back again. Then, about one half of this amount is attributed to a delay created for the one way passage of a signal to the client.

- [00298] The host and client both drive the line to a logic-zero level during both guard times to keep the MDDI_DATA lines in a defined state. The enable and disable times of the host and client during both guard times are such that the MDDI_Data signals are at a valid low level for any valid round-trip delay time.

25. Forward Link Skew Calibration Packet

- [00299] The Forward Link Skew Calibration Packet allows a client or display to calibrate itself for differences in the propagation delay of the MDDI_Data signals with respect to the MDDI_Stb signal. Without delay skew compensation, the maximum data rate is generally limited to account for potential worst-case variation in these delays. Generally, this packet is only sent when the forward link data rate is configured to a rate of around 50 Mbps or lower. After sending this packet to calibrate the display, the data rate may be stepped up above 50 Mbps. If the data rate is set too high during the skew calibration process, the display might synchronize to an alias of the bit period which could cause the delay skew compensation setting to be off by more than one bit time, resulting in erroneous data clocking. The highest data rate type of interface or greatest possible Interface Type is selected prior to sending the Forward Link Skew Calibration Packet so that all existing data bits are calibrated.

- [00300] One embodiment of the format of a Forward Link Skew Calibration Packet is shown in FIG. 56. As shown in FIG. 56, this type of packet is structured to have Packet Length (2 bytes), Packet Type, hClient ID, Parameter CRC, All Zero 1, Calibration Data Sequence, and All Zero 2 fields. This type of packet is generally identified as a Type 83 packet in the type field, and in one embodiment has a pre-selected length of 519.

Virtual Control Panel

- [00301] The use of a VCP allows a host to set certain user controls in a client. By allowing these parameters to be adjusted by the host, the user interface in the client can be simplified because screens that allow a user to adjust parameters such as audio

volume or display brightness can be generated by host software rather than by one or more microprocessors in the client. The host has the ability to read the parameter settings in the client and to determine the range of valid values for each control. The client generally has the capability to report back to the host which control parameters can be adjusted.

- [00302] The control codes (VCP Codes) and associated data values generally specified, are utilized to specify controls and settings in the client. The VCP Codes in the MDDI specification are expanded to 16 bits to preserve proper data field alignment in the packet definitions, and in the future to support supplementary values that are unique to this interface or future enhancements.

26. Request VCP Feature Packet

- [00303] The Request VCP Feature Packet provides a means, mechanism, or method for the host to request the current setting of a specific control parameter or all valid control parameters. Generally, a client responds to a VCP Packet with the appropriate information in a VCP Feature Reply Packet. In one embodiment, the client indicates an ability to support the Request VCP Feature Packet using bit 13 of the Client Feature Capability Indicators field of the Client Capability Packet.
- [00304] The format of the Request VCP Feature Packet in one embodiment is shown in FIG. 69. As seen in FIG. 69, this type of packet is structured to have Packet Length, Packet Type, hClient ID, Monitor Control Command Set (MCCS) VCP code, and CRC fields. This type of packet is generally identified in one embodiment as a Type 128, which is indicated in the 2 byte type field. The packet length, which specifies the total number of bytes in the packet not including the packet length field, is typically fixed for this type of packet at a length of 8 bytes.
- [00305] The hClient ID field is reserved for use as a Client ID in future implementations and is typically set to zero. The MCCS VCP Code field comprises 2 bytes of information that specifies the MCCS VCP Control Code Parameter. A value in the range of 0 to 255 causes a VCP Feature Reply Packet to be returned with a single item in the VCP Feature Reply List corresponding to the specified MCCS code. An MCCS VCP Code of 65535 (0xffff) requests a VCP Feature Reply Packet with a VCP Feature Reply List containing a Feature Reply List Item for each control supported by the client.

The values of 256 through 65534, for this field are reserved for future use and presently not in use.

27. *VCP Feature Reply Packet*

- [00306] The VCP Feature Reply Packet provides a means, mechanism, or method for a client to respond to a host request with the current setting of a specific control parameter or all valid control parameters. Generally, a client sends the VCP Feature Reply Packet in response to a Request VCP Feature Packet. This packet is useful to determine the current setting of a specific parameter, to determine the valid range for a specific control, to determine if a specific control is supported by the client, or to determine the set of controls that are supported by the client. If a Request VCP Feature is sent that references a specific control that is not implemented in the client then a VCP Feature Reply Packet is returned with a single VCP Feature Reply List item corresponding to the unimplemented control that contains the appropriate error code. In one embodiment, the client indicates an ability to support the VCP Feature Reply Packet using bit 13 of the Client Feature Capability field of the Client Capability Packet.
- [00307] The format of the VCP Feature Reply Packet in one embodiment is shown in FIG. 70. As seen in FIG. 70, this type of packet is structured to have Packet Length, Packet Type, cClient ID, MCCS Version, Reply Sequence Number, VCP Feature Reply List, and CRC fields. This type of packet is generally identified in one embodiment as a Type 129, as indicated in the 2 byte type field.
- [00308] The cClient ID field contains information reserved for a Client ID. This field is reserved for future use and is generally set to zero. MCCS Version field contains 2 bytes of information that specifies the Version of the VESA MCCS Specification implemented by the client.
- [00309] The 2 byte Reply Sequence Number field contains information or data that specifies the sequence number of the VCP Feature Reply Packets returned by the client. The client returns one or more VCP Feature Reply Packets in response to a Request VCP Feature Packet with an MCCS Control Code value of 65535. The client may spread or transfer the feature reply list over multiple VCP Feature Reply Packets. In this case, the client should assign a sequence number or identifier to each successive packet, and the sequence numbers of the VCP Feature Reply Packets sent in response to a single Request VCP Feature Packet typically starts at zero and increments by one.

The last VCP Feature Reply List Item in the last VCP Feature Reply Packet should contain an MCCS VCP Control Code value equal to 0xffff to identify that the packet is the last one and contains the highest sequence number of the group of packets returned. If only one VCP Feature Reply Packet is sent in response to a Request VCP Feature Packet then the Reply Sequence Number in that single packet is generally set at zero and the VCP Feature Reply List contains a list item having an MCCS VCP Code in the VCP Feature Reply List Item equal to 0xffff. The Maximum Value and Present Value fields (FIG. 71) in the VCP Feature Reply List Item packet are set to zero when the MCCS VCP Control Code is equal to 0xffff.

[00310] The Number of Features in List field contains 2 bytes that specify the number of VCP Feature Reply List Items that are in the VCP Feature Reply List in this packet, while the VCP Feature Reply List field is a group of bytes that contain one or more VCP Feature Reply List Items. The format of a single VCP Feature Reply List Item in one embodiment is shown in FIG. 71.

[00311] As shown in FIG. 71, each VCP Feature Reply List Item is 12 bytes in length, and comprises the MCCS VCP Code, Result Code, Maximum Value, and Present Value fields. The 2-byte MCCS VCP Code field contains data or information that specifies the MCCS VCP Control Code Parameter associated with this list item. Only the Control Code values defined in the VESA MCCS Specification version 2 and later are considered as valid for this embodiment. The 2-byte Result Code field contains information that specifies an error code related to the request for information regarding the specified MCCS VCP Control. A value of '0' in this field means there is no error, while a value of '1' means the specified control is not implemented in the client. Further values for this field of 2 through 65535 are currently reserved for future use and implementation of other application contemplated by the art, but are not to be used for now.

[00312] The 4-byte Maximum Value field specifies the largest possible value to which the specified MCCS Control can be set. If the requested control is not implemented in the client this value is set to zero. If the value returned is less than 32 bits (4 bytes) in length, then the value is cast into a 32-bit integer leaving the most significant (unused) bytes set to zero. The 4-byte Present Value field contains information that specifies the present value of the specified MCCS VCP Continuous (C) or Non-Continuous (NC) control. If the requested control is not implemented in the client or if the control is

implemented but is a Table (T) data type, then this value is set to zero. If the value returned is less than 32 bits (4 bytes) in length per the VESA MCCS specification then the value is cast into a 32-bit integer leaving the most significant (unused) bytes set to zero. If the specified MCCS VCP code corresponds to a non-continuous control or table data type, then the Maximum Value field is set or selected to be zero.

28. *Set VCP Feature Packet*

- [00313] The Set VCP Feature Packet provides a means, mechanism, or method for a host to set VCP control values for both continuous and non-continuous controls in a client. In one embodiment, the client indicates the ability to support the Set VCP Feature Packet using bit 13 of the Client Feature Capability field of the Client Capability Packet.
- [00314] The format of the Set VCP Feature Packet in one embodiment is shown in FIG. 72. As seen in FIG. 72, this type of packet is structured to have Packet Length, Packet Type, hClient ID, MCCS VCP Code, Number of Values in List, Control Value List, and CRC fields. This type of packet is generally identified as a Type 130, as indicated in the 2 byte type field, is 20 bytes long exclusive of the Packet Length field.
- [00315] The hClient ID field again uses a 2-byte value to specify or act as a Client ID. This field is reserved for future use and is currently set to zero. The MCCS VCP Code field uses 2 bytes of information or values to specify the MCCS VCP Control Code Parameter to be adjusted. The 2-byte Number of Values in List Field contains information or values that specify the number of 16-bit values that exist in the Control Value List. The Control Value List will usually contain one item unless the MCCS Control Code relates to a table in the client. In the case of non-table-related controls, The Control Value List will contain a value that specifies the new value to be written to the control parameter specified by the MCCS VCP Code field. For table-related controls the format of the data in the Control Value List is specified by the parameter description of the specified MCCS VCP Code. If the list contains values that are larger than one byte, then the least-significant byte is transmitted first, consistent with the method defined elsewhere. Finally, the 2-byte CRC field contains a 16-bit CRC of all bytes in the packet including the Packet Length.

29. *Request Valid Parameter Packet*

[00316] The Request Valid Parameter Packet is used as a means or structure useful to request that a client return a Valid Parameter Reply Packet containing a list of parameters supported by the specified NC or Table (T) control. This packet should only specify non-continuous controls or controls that relate to a table in the client, and not specify a MCCS VCP Code value of 65535 (0xffff) to specify all controls. If a non-supported or invalid MCCS VCP Code is specified then an appropriate error value is returned in the Valid Parameter Reply Packet. In one embodiment, the client indicates an ability to support the Request Valid Parameter Packet using bit 13 of the Client Feature Capability field of the Display Capability Packet.

[00317] The format of the Request Valid Parameter Packet in one embodiment is shown in FIG. 73. As seen in FIG. 73, this type of packet is structured to have Packet Length, Packet Type, hClient ID, MCCS VCP Code, and CRC fields. This type of packet is generally identified in one embodiment as a Type 131, as indicated in the 2 byte type field.

[00318] The packet length, as indicated in the 2-bytes Packet Length Field is generally set to have a total number of bytes in the packet, not including the packet length field of 8. The hClient ID again specifies the Client ID, but is currently reserved for future use, as would be apparent to one skilled in the art, and is set to zero. The 2-byte MCCS VCP Code Field contains a value that specifies the non-continuous MCCS VCP Control Code Parameter to be queried. The value in this field should correspond to a non-continuous control that is implemented in the client. The values 256 through 65535 (0xffff) are typically reserved or considered as invalid, and are considered as an unimplemented control in the error response.

30. *Valid Parameter Reply Packet*

[00319] A Valid Parameter Reply Packet is sent in response to a Request Valid Parameter Packet. It is used as a means, method, or structure to identify the valid settings for a non-continuous MCCS VCP control or a control that returns the contents of a table. If the control relates to a table in the client, then the VCP Parameter Reply List simply contains the specific list of sequential table values that were requested. If the contents of the table cannot fit into a single Valid Parameter Reply Packet then multiple packets with sequential Reply Sequence Numbers can be sent by the client. In

one embodiment, a client indicates an ability to support a Valid Parameter Reply Packet using bit 13 of the Client Feature Capability field of the Client Capability Packet.

[00320] A host may request the contents of a table in the following manner: the host sends a Set VCP Feature Packet containing the necessary or desired parameters such as read/write parameter, Look-Up Table (LUT) offset, and RGB selection; then a Request Valid Parameter Packet that specifies the desired control is sent by the host; then the client returns one or more Valid Parameter Reply Packets containing the table data. This sequence of operations performs a similar function as the table reading functions described in the MCCS operation model.

[00321] If a specific client parameter is not supported by the client then in one embodiment the corresponding field of this packet will contain a value of 255. For parameters that are used in the client, the corresponding field should contain a value of the parameter in the client.

[00322] The format of the Valid Parameter Reply Packet for one embodiment is shown in FIG. 74. As seen in FIG. 74, this type of packet is structured to have Packet Length, Packet Type, cClient ID, MCCS VCP Code, Response Code, Reply Sequence Number, Number Values in List, VCP Parameter Reply List, and CRC fields. This type of packet is generally identified for one embodiment as a Type 132, as indicated in the 2 byte type field.

[00323] The cClient ID field is reserved for the future Client ID, as is known from the above discussions, while the 3-byte MCCS VCP Code Packet contains a value that specifies a non-continuous MCCS VCP Control Code Parameter that is described by this packet. If an invalid MCCS VCP Control Code is specified by a Request Valid Parameter Packet, then the same invalid parameter value will be specified in this field with the appropriate value in the Response Code field. If the MCCS Control Code is invalid then the VCP Parameter Reply List will have zero length.

[00324] The Response Code field contains 2 bytes of information or values that specify the nature of the response related to the request for information regarding the specified MCCS VCP Control. If the value in this field is equal to 0, then no error is considered as being present for this data type, and the last Valid Parameter Reply Packet in the sequence is sent, it having the highest Reply Sequence Number. If the value in this field is equal to 1, then no error is considered as being present, but other Valid Parameter Reply Packets will be sent that have higher sequence numbers. If the value in this field

is equal to 2, then the specified control is not considered as being implemented in the client. If the value in this field is equal to 3, then the specified control is not a non-continuous control (it is a continuous control that always has a valid set of all values from zero to its maximum value). Values for this field equal to 4 through 65535 are reserved for future use and generally not to be used.

[00325] The 2-byte Reply Sequence Number field specifies the sequence number of the Valid Parameter Reply Packets returned by the client. The client returns one or more Valid Parameter Reply Packets in response to a Request Valid Parameter Packet. The client may spread the VCP Parameter Reply List over multiple Valid Parameter Reply Packets. In this latter case, the client will assign a sequence number to each successive packet, and set the Response Code to 1 in all but the last packet in the sequence. The last Valid Parameter Reply Packet in the sequence will have the highest Reply Sequence Number and the Response Code contains a value of 0.

[00326] The 2-byte Number of Values in List field specifies the number of 16-bit values that exist in the VCP Parameter Reply List. If the Response Code is not equal to zero then the Number of Values in List parameter is zero. The VCP Parameter Reply List field contains a list of 0 to 32760 2-byte values that indicate the set of valid values for the non-continuous control specified by the MCCS Control Code field. The definitions of the non-continuous control codes are specified in the VESA MCCS Specification. Finally, in this embodiment, the CRC field contains a 16-bit CRC of all bytes in the packet including the Packet Length.

Scaled Video Stream Images

[00327] The MDDI or protocol mechanism, structure, means, or method provides support for scaled video stream images that allow the host to send an image to the client that is scaled larger or smaller than the original image, and the scaled image is copied to a main image buffer. An overview of the Scaled Video Stream functionality and associated protocol support is provided elsewhere. An ability to support scaled video streams is defined by or within the Scaled Video Stream Capability Packet, which is sent in response to a Request Specific Status Packet.

[00328] The header of the Scaled Video Stream packet discussed below is slightly different from the simpler Video Stream Packet whose header contains the entire context necessary to display the image. The Scaled Video Stream packet uses a setup

packet to define the parameters of the source and destination window size and position, and a separate Scaled Video Stream Packet to transmit the pixel data. A client allocates internal storage associated with each stream to store the stream parameters from the setup packet and part of the pixel data associated with each stream. The amount of storage required for each stream will vary depending on the size of the source and destination images and the values specified in the setup packet. For this reason the protocol is designed to allow the client to implement dynamic memory allocation for the assignment of storage associated with each scaled video stream.

- [00329] It is useful to send a video stream to a display having a size native to the program source and have the display scale and position the image in a manner appropriate for the specific end application. The implementation for real-time scaling of multiple video images is sufficiently complex to make support of this feature optional in the client.

31. Scaled Video Stream Capability Packet

- [00330] The Scaled Video Stream Capability Packet defines the characteristics of the scaled video stream source image in or used by a client. The format of the Scaled Video Stream Capability Packet is shown generally in FIG. 75. As seen in FIG. 75, in one embodiment, a Scaled Video Stream Capability Packet is structured to have Packet Length, Packet Type, cClient ID, Max Number of Streams, Source Max X Size, Source Max Y size, RGB Capability, Monochrome Capability, Reserved 1, Y Cr Cb Capability, Reserved 2, and CRC fields. The packet length, in one embodiment, is selected to be a fixed 20 bytes, as shown in the length field, including the 2-byte cClient ID field, which is reserved for use for a Client ID, otherwise set to zero, and the CRC field. In one embodiment, the client indicates an ability to support the Scaled Video Stream Capability Packet using a parameter value of 143 in the Valid Parameter Reply List of the Valid Status Reply List Packet.

- [00331] The 2-byte Maximum Number of Streams field contains a value to identify the maximum number of simultaneous scaled video streams that may be allocated at one time. In one embodiment, a client should deny a request to allocate a scaled video stream if the maximum number of scaled video streams is already allocated. If less than the maximum number of scaled video streams are allocated the client may also deny an allocation request based on other resource limitations in the client.

- [00332] The Source Maximum X Size and Y size fields (2 bytes) specify values for the maximum width and height, respectively, of the scaled video stream source image expressed as a number of pixels.
- [00333] The RGB Capability field uses values to specify the number of bits of resolution that can be displayed in RGB format. If the scaled video stream cannot use the RGB format then this value is set equal to zero. The RGB Capability word is composed of three separate unsigned values with: Bits 3 through 0 defining a maximum number of bits of blue (the blue intensity) in each pixel, Bits 7 through 4 defining the maximum number of bits of green (the green intensity) in each pixel, and Bits 11 through 8 defining the maximum number of bits of red (the red intensity) in each pixel, while Bits 15 through 12 are reserved for future use in future capability definitions, and are generally set to zero.
- [00334] The 1-byte Monochrome Capability field contains a value that specifies the number of bits of resolution that can be displayed in monochrome format. If the scaled video stream cannot use the monochrome format then this value is set to zero. Bits 7 through 4 are reserved for future use and should, therefore, be set to zero ('0') for current applications, although this may change over time, as will be appreciated by those skilled in the art. Bits 3 through 0 define the maximum number of bits of grayscale that can exist in each pixel. These four bits make it possible to specify that each pixel consists of 1 to 15 bits. If the value is zero, then the monochrome format is not supported by the scaled video stream.
- [00335] The Reserved 1 field (here 1 byte) is reserved for future use in providing values related to the Scaled Video Stream Packet information or data. Therefore, currently, all bits in this field are set to a logic '0'. One purpose of this field is to cause all subsequent 2-byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.
- [00336] The 2-byte Y Cb Cr Capability field contains values that specify the number of bits of resolution that can be displayed in Y Cb Cr format. If the scaled video stream cannot use the Y Cb Cr format then this value is zero. The Y Cb Cr Capability word is composed of three separate unsigned values with: Bits 3 through 0 defining the maximum number of bits that specify the Cr sample; Bits 7 through 4 defining the maximum number of bits that specify the Cb sample; Bits 11 through 8 defining the

maximum number of bits specify the Y sample; and with Bits 15 through 12 being reserved for future use and is generally set to zero.

[00337] The 1-byte Capability Bits field contains a set of flags that specify capabilities associated with the scaled video stream. The flags are defined as follows: Bit 0 covers Pixel data in the Scaled Video Stream Packet can be in a packed format. An example of packed and byte-aligned pixel data is shown earlier in FIG. 13. Bit 1 is reserved for future use and is generally set to zero; Bit 2 is also reserved for future use and is set to zero; Bit 3 covers scaled video streams that can be specified in the color map data format. The same color map table is used for the scaled video streams as is used for the main image buffer and the alpha-cursor image planes. The color map is configured using the Color Map Packet described elsewhere; and Bits 7 through 4 are reserved for future use and are generally set to be zero.

[00338] The Reserved 2 field (here 1 byte) is reserved for future use in providing values related to the Scaled Video Stream Packet information or data. Therefore, currently, all bits in this field are set to a logic '0'. One purpose of this field is to cause all subsequent 2-byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.

32. Scaled Video Stream Setup Packet

[00339] The Scaled Video Stream Setup Packet provides a means, structure, or method used to define the parameters of the scaled video stream and the client uses the information to allocate internal storage for buffering and scaling of the image. A stream may be de-allocated by sending this packet with the X Image Size and Y Image Size fields equal to zero. Scaled video streams that have been de-allocated may be reallocated later with the same or different stream parameters. In one embodiment a client indicates an ability to support the Scaled Video Stream Setup Packet using a parameter value of 143 in the Valid Parameter Reply List of the Valid Status Reply List Packet, and by using a non-zero value in the Maximum Number of Streams field of the Scaled Video Stream Capability Packet.

[00340] The format of the Scaled Video Stream Setup Packet is shown generally in FIG. 76. As seen in FIG. 76, in one embodiment, a Scaled Video Stream Setup Packet is structured to have Packet Length, Packet Type, hClient, Stream ID, Video Data

Format Descriptor, Pixel Data Attributes, X Left Edge, Y Top Edge, X Right Edge, Y Bottom Edge, X Image Size, Y Image Size, and CRC fields.

[00341] The 2-byte Packet Length field specifies the total number of bytes in the packet not including the packet length field. In one embodiment, this packet length is fixed at 24. The 2-byte Packet Type field employs a value of 136 to identify the packet as a Scaled Video Stream Setup Packet. The 2-byte hClient ID field is reserved for future use as a Client ID, and is generally set to an all bits at logic-zero value for the moment, or until a protocol user determines what ID values are to be used, as would be known.

[00342] The Stream ID field uses 2 bytes to specify a unique identifier for the Stream ID. This value is assigned by the host and ranges in value from zero to the maximum Stream ID value specified in the Client Capability Packet. The host must manage the use of Stream ID values carefully to ensure that each active stream is assigned a unique value, and that streams that are no longer active are de-allocated or reassigned.

[00343] In one embodiment, the Video Data Format Descriptor field uses 2 bytes to specify the format of each pixel in the Pixel Data in the present stream in the present packet. The pixel data format should comply with at least one of the valid formats for a alpha-cursor image plane as might be defined in an Alpha-Cursor Image Capability Packet, or other pre-defined image pattern, as will generally be defined within the other packets discussed above. The Video Data Format Descriptor defines the pixel format for the current packet only and does not imply that a constant format will continue to be used for the lifetime of a particular video stream. Fig. 12 illustrates an embodiment of how the Video Data Format Descriptor is coded, and as discussed above for other packets.

[00344] For example, as seen in FIGs. 12A through 12D, and for use in one embodiment, when bits [15:13] are equal to '000', then the video data consists of an array of monochrome pixels where the number of bits per pixel is defined by bits 3 through 0 of the Video Data Format Descriptor word. Bits 11 through 4 are generally reserved for future use or applications and are set to zero in this situation. When bits [15:13] are instead equal to the values '001', then the video data consists of an array of color pixels that each specify a color through a color map (palette). In this situation, bits 5 through 0 of the Video Data Format Descriptor word define the number of bits per pixel, and bits 11 through 6 are generally reserved for future use or applications and set equal to zero. When bits [15:13] are instead equal to the values '010', then the video data consists of an

array of color pixels where the number of bits per pixel of red is defined by bits 11 through 8, the number of bits per pixel of green is defined by bits 7 through 4, and the number of bits per pixel of blue is defined by bits 3 through 0. In this situation, the total number of bits in each pixel is the sum of the number of bits used for red, green, and blue.

[00345] However, when bits [15:13] are instead equal to the values or string '011', as shown in FIG. 12D, then the video data consists of an array of video data in 4:2:2 YCbCr format with luminance and chrominance information, where the number of bits per pixel of luminance (Y) is defined by bits 11 through 8, the number of bits of the Cb component is defined by bits 7 through 4, and the number of bits of the Cr component is defined by bits 3 through 0. The total number of bits in each pixel is the sum of the number of bits used for red, green, and blue. The Cb and Cr components are sent at half the rate as Y. In addition, the video samples in the Pixel Data portion of this packet are organized as follows: Cbn, Yn, Cm, Yn+1, Cbn+2, Yn+2, Cm+2, Yn+3, ... where Cbn and Cm are associated with Yn and Yn+1, and Cbn+2 and Cm+2 are associated with Yn+2 and Yn+3, and so on. Yn, Yn+1, Yn+2 and Yn+3 are luminance values of four consecutive pixels in a single row from left to right

[00346] For all four formats discussed above. Bit 12, which is designated as "P" in the figures, specifies whether or not the Pixel Data samples are packed, or byte-aligned pixel data. A value of '0' in this field indicates that each pixel in the Pixel Data field is byte-aligned with an MDDI byte boundary. A value of '1' indicates that each pixel and each color within each pixel in the Pixel Data is packed up against the previous pixel or color within a pixel leaving no unused bits.

[00347] In one embodiment, a 2-byte Pixel Data Attributes field has values that are interpreted as follows with Bits 1 and 0 being reserved for future use, generally set to logic-zero for now, and Bit 2 indicates whether or not the Pixel Data is in interlace format. When Bit 2 is 0, then the Pixel Data is in the standard progressive format. The row number (pixel Y coordinate) is incremented by 1 when advancing from one row to the next. When Bit 2 is 1, then the Pixel Data is in interlace format. The row number (pixel Y coordinate) is incremented by 2 when advancing from one row to the next.

[00348] In one embodiment, Bit 3 indicates whether or not the Pixel Data is in alternate pixel format. This is similar to the standard interlace mode enabled by bit 2, but with the interlacing being vertical rather than horizontal. When Bit 3 is 0, the Pixel Data is

generated or placed in the standard progressive format. The column number (pixel X coordinate) is incremented by 1 as each successive pixel is received. When Bit 3 is 1, then the Pixel Data is generated or placed in alternate pixel format. The column number (pixel X coordinate) is incremented by 2 as each pixel is received.

[00349] Bits 4 through 15 are also reserved for future use, and are generally be set to logic-zero level or values for current applications or designs.

33. *Scaled Video Stream Acknowledgement Packet*

[00350] The Scaled Video Stream Acknowledgement Packet allows a client to acknowledge the receipt of a Scaled Video Stream Setup Packet. The client can indicate an ability to support the Scaled Video Stream Acknowledgement Packet via a parameter value of 143 in the Valid Parameter Reply List of the Valid Status Reply List Packet and via a non-zero value in the Maximum Number of Streams field of the Scaled Video Stream Capability Packet.

[00351] The format of the Scaled Video Stream Acknowledgement Packet is shown generally in FIG. 77. As seen in FIG. 77, in one embodiment, a Scaled Video Stream Acknowledgement Packet is structured to have Packet Length, Packet Type, cClient, Stream ID, ACK Code, and CRC fields. The 2-byte Packet Length field is used to specify the total number of bytes, excluding the packet length field, with a value of 10 for this packet type, while a Packet Type of 137 identifies a packet as a Scaled Video Stream Acknowledgement Packet.

[00352] The 2-byte cClient ID field is reserved for future use for the Client ID, and is generally set to zero. The 2-byte Stream ID field specifies a unique identifier for the Stream ID. This is the same value assigned by the host in the Scaled Video Stream Setup Packet.

[00353] The 2-byte Ack Code field provides values containing a code that describes the outcome of an attempt to update the specified scaled video stream. In one embodiment, the codes are defined as follows:

- 0 – The stream allocation attempt was successful.
- 1 – the stream de-allocation attempt was successful.
- 2 – invalid attempt to allocate a stream ID that has already been allocated.
- 3 – invalid attempt to de-allocate a stream ID that is already de-allocated.
- 4 – the client does not support scaled video streams

- 5 – the stream parameters are inconsistent with the capability of the client.
- 6 – stream ID value larger than the maximum value allowed by the client.
- 7 – insufficient resources available in the client to allocate the specified stream.

[00354] The 2-byte CRC field contains the CRC of all bytes in the packet including the Packet Length.

34. Scaled Video Stream Packet

[00355] The Scaled Video Stream Packet is used to transmit the pixel data associated with a specific scaled video stream. The size of the region reference by this packet is defined by the Scaled Video Stream Setup Packet. The client can indicate an ability to support the Scaled Video Stream Packet using a parameter value of 143 in the Valid Parameter Reply List of the Valid Status Reply List Packet and using a successful scaled video stream allocation response in the Ack Code field of the Scaled Video Stream Acknowledgement Packet.

[00356] The format of one embodiment of the Scaled Video Stream Packet is shown generally in FIG. 78. As seen in FIG. 78, a Scaled Video Stream Packet is structured to have Packet Length, Packet Type, hClient ID, Stream ID, Pixel data Attributes, Pixel Count, Parameter CRC, Pixel Data, and Pixel Data CRC fields. The 2-byte Packet Type field uses a value of 18 to identify a packet as a Scaled Video Stream Packet. The hClient ID field is reserved for the Client ID, and generally set to zero. As before, the 2-byte Stream ID field specifies a unique identifier for the Stream ID. This value is specified by the host in the Scaled Video Stream Setup Packet and confirmed in the Scaled Video Stream Acknowledgement Packet. In one embodiment, the 2-byte Pixel Data Attributes field has values that specify pixel data routing and display updating or buffer locations. These values are, in one embodiment,

[00357] In one embodiment, the 2-byte Pixel Data Attributes field has values that specify pixel data routing and display updating or buffer locations. These values are, in one embodiment, interpreted as follows: with Bits 1 and 0 selecting the display where the pixel data is to be routed. For bit values of '11' or '00' pixel data is displayed to or for both eyes, for bit values '10', pixel data is routed only to the left eye, and for bit values '01', and pixel data is routed only to the right eye.

[00358] Bits 7 and 6 are the Display Update Bits that specify the frame buffer where the pixel data is to be written. The effects of the Frame Update Bits are described in more

detail elsewhere. When Bits [7:6] are '01', the Pixel data is written to the offline image buffer. When Bits [7:6] are '00', the Pixel data is written to the image buffer used to refresh the display. When Bits [7:6] are '11', the Pixel data is written to all image buffers. If Bits [7:6] are '10', this is treated as an invalid value. These bits are currently reserved for future use. In this situation, Pixel data would be ignored and not written to any of the image buffers. Bits 2 through 5 and 8 through 15 are reserved for future use and are generally be set to logic-zero level or values.

[00359] The 2-byte Pixel Count field specifies the number of pixels in the Pixel Data field below. The 2-byte Parameter CRC field has the CRC of all bytes from the Packet Length to the Pixel Count. If this CRC fails to check then the entire packet is discarded. The 2-byte Pixel Data field contains the raw video information that is to be scaled and then displayed. Data is formatted in the manner described by the Video Data Format Descriptor field. The data is transmitted a row at a time as defined previously.

[00360] The 2-byte Pixel Data CRC field contains a CRC of only the Pixel Data. If this CRC fails to check then the Pixel Data can still be used but the CRC error count is incremented.

35. *Request Specific Status Packet*

[00361] The Request Specific Status Packet provides a means, mechanism, or method for a host to request that the client send a capability or status packet back to the host as specified in this packet. The client returns the packet of the specified type in the next Reverse Link Encapsulation Packet. The client will generally set bit 17 in the Client Feature Capability field of the Client Capability Packet if the client has the capability to respond to the Request Specific Status Packet. A convenient method for the host to use to determine all of the types of status packets that a client can return or transfer is to use the Valid Status Reply List Packet described elsewhere. The client can indicate an ability to respond with the Valid Status Reply List Packet using bit 21 of Client Feature Capability field of the Client Capability Packet.

[00362] The format of one embodiment of a Request Specific Status Packet is shown generally in FIG. 79. As seen in FIG. 79, a Request Specific Status Packet is structured to have Packet Length, Packet Type, hClient ID, Status Packet ID, and CRC fields. Packet Length field specifies the total number of bytes in the packet not including the packet length field, and is generally fixed at a value of 10 for this packet type. A Packet

Type of 138 identifies the packet as a Request Specific Status Packet. The hClient ID field (2 bytes) is reserved for future use for a Client ID, and is set to zero for now, while a 2-byte Status Packet ID field specifies the type of capability or status packet that the client is going to send to the host. Typical packets types are:

66 – Client Capability Packet is sent by the client.

133 – Alpha-Cursor Image Capability Packet is sent by the client.

139 – Valid Status Reply List Packet is sent that identifies the exact types of capability and status packets that the client can send.

141 – Personal Client Capability Packet is sent by the client.

142 – Client Error Report Packet is sent by the client.

143 – Scaled Video Stream Capability Packet is sent by the client.

144 – Client Identification Packet is sent by the client.

[00363] Packet Types 56 through 63 can be used for manufacturer-specific capability and status identifiers.

[00364] The CRC field again contains a CRC of all bytes in the packet including the Packet Length.

36. Valid Status Reply List Packet

[00365] The Valid Status Reply List Packet provides the host with a structure, means, or method to have a list of status and capability packets that the client has the capability to return. A client can indicate an ability to support the Valid Status Reply List Packet using bit 21 of Client Feature Capability field of the Client Capability Packet.

[00366] The format of one embodiment of a Valid Status Reply List Packet is shown generally in FIG. 80. As seen in FIG. 80, a Valid Status Reply List Packet is structured to have Packet Length, Packet Type, cClient ID, Number of Values in List, Valid Parameter Reply List, and CRC fields. The packet length for this type of packet is generally fixed at a value of 10, and a type value of 139 identifies the packet as a Valid Status Reply Packet. The cClient ID field is reserved for future use as the Client ID, and is generally be set to zero. The 2- byte Number of Values in List field specifies the number of items in the following Valid Parameter Reply List.

[00367] The Valid Parameter Reply List field contains a list of 2-byte parameters that specify the types of capability or status packets that the client can send to the host. If the client has indicated that it can respond to the Request Specific Status Packet (using

bit 21 of the Client Feature Capability field the in the Client Capability Packet) then it is capable of sending at least the Client Capability Packet (Packet Type = 66) and the Valid Status Reply List Packet (Packet Type = 139). The Packet Types that can be sent by the client and may be included in this list, along with their respective assignments for purposes of the one embodiment, are:

- 66 – Client Capability Packet.
- 133 – Alpha-Cursor Image Capability Packet.
- 139 – Valid Status Reply List Packet, that identifies the exact types of capability and status packets that the client can send.
- 141 – Personal Display Capability Packet.
- 142 – Client Error Report Packet.
- 143 – Scaled Video Stream Capability Packet.
- 144 – Client Identification Packet.
- 145 – Alternate Display Capability Packet.

[00368] Packet Types 56 through 63 can be used for manufacturer-specific capability and status identifiers.

[00369] The CRC field contains a CRC of all bytes in the packet including the Packet Length.

37. *Personal Display Capability Packet*

[00370] The Personal Display Capability Packet provides a set of parameters that describe the capabilities of a personal display device, such as a head-mounted display or display glasses. This enables the host to customize the display information according to the specific capabilities of a client. A client, on the other hand, indicates an ability to send the Personal Display Capability Packet by using a corresponding parameter in the Valid Parameter Reply List of the Valid Status Reply List Packet.

[00371] The format of one embodiment of a Personal Display Capability Packet is shown generally in FIG. 81. As seen in FIG. 81, a Personal Display Capability Packet is structured to have Packet Length, Packet Type, cClient ID, Sub-Pixel Layout, Pixel Shape, Horizontal Field of View, Vertical Field of View, Visual Axis Crossing, Lft./Rt. Image, See Through, Maximum Brightness, Optical Capability, Minimum IPD, Maximum IPD, Points of Ifeld of Curvature List and CRC fields. In one embodiment, the Packet Length field value is fixed at 68. A Packet Type value of 141 identifies a

packet as a Personal Display Capability Packet. The cClient ID field is reserved for future use and is generally set to zero for now.

[00372] The Sub-Pixel Layout field specifies the physical layout of a sub-pixel from top to bottom and left to right, using values of: 0 to indicate that a sub-pixel layout is not defined; 1 to indicate red, green, blue stripe; 2 to indicate blue, green, red stripe; 3 to indicate a quad-pixel, having a 2-by-2 sub-pixel arrangement of red at the top left, blue at the bottom right, and two green sub-pixels, one at the bottom left and the other at the top right; 4 to indicate a quad-pixel, with a 2-by-2 sub-pixel arrangement of red at the bottom left, blue at the top right, and two green sub-pixels, one at the top left and the other at the bottom right; 5 to indicate a Delta (Triad); 6 to indicate a mosaic with red, green, and blue overlaid (e.g. LCOS display with field-sequential color); and with values 7 through 255 being generally reserved for future use.

[00373] The Pixel Shape field specifies the shape of each pixel that is composed of a specific configuration sub-pixels, using a value of: 0 to indicate that a sub-pixel shape is not defined; 1 to indicate round; 2 to indicate square; 3 to indicate rectangular; 4 to indicate oval; 5 to indicate elliptical; and with the values 6 through 255 being reserved for future use in indicating desired shapes, as can be appreciated by one skilled in the art.

[00374] A 1-byte Horizontal Field of View (HFOV) field specifies the horizontal field of view in 0.5 degree increments (e.g. if the HFOV is 30 degrees, this value is 60). If this value is zero then the HFOV is not specified.

[00375] A 1-byte Vertical Field of View (VFOV) field specifies the vertical field of view in 0.5 degree increments (e.g. if the VFOV is 30 degrees, this value is 60). If this value is zero then the VFOV is not specified.

[00376] A 1-byte Visual Axis Crossing field specifies the visual axis crossing in 0.01 diopter (1/m) increments (e.g. if the visual axis crossing is 2.22 meters, this value is 45). If this value is zero then the Visual Axis Crossing is not specified.

[00377] A 1-byte Left/Right Image Overlap field specifies the percentage of overlap of the left and right image. The allowable range of the image overlap in percent is 1 to 100. Values of 101 to 255 are invalid and are generally not to be used. If this value is zero then the image overlap is not specified.

[00378] A 1-byte See Through field specifies the see-through percentage of image. The allowable range of see-through in percent is 0 to 100. Values of 101 to 254 are invalid

and are not to be used. If this value is 255 then the see-through percentage is not specified. A1-byte Maximum Brightness field specifies the maximum brightness in increments of 20 nits (e.g. if the maximum brightness is 100 nits, this value is 5). If this value is zero then the maximum brightness is not specified.

[00379] A 2-byte Optical Capability Flags field contains various fields that specify optical capabilities of the display. These bit values are generally assigned according to:

[00380] Bits 15 through 5 are reserved for future use and are generally set to a logic-zero state.

[00381] Bit 4 selects Eye Glass Focus Adjustment, with a value of '0' meaning the display has no eye glass focus adjustment, and a value of '1' meaning the display has an eye glass focus adjustment.

[00382] Bits 3 through 2 select a Binocular Function according to: a value of 0 means the display is binocular and can display 2-dimensional (2D) images only; 1 means the display is binocular and can display 3-dimensional (3D) images; 2 means the display is monocular, and 3 is reserved for future use.

[00383] Bits 1 through 0 select Left-Right Field Curvature Symmetry, with a value of 0 meaning Field curvature not defined. If this field is zero then all field curvature values from A1 through E5 are set to zero except for point C3, which specifies a focal distance of the display or is to be set to zero to indicate the focal distance is not specified. A value of 1 means Left and Right displays have the same symmetry; 2 means Left and right displays are mirrored on the vertical axis (column C); and 3 is reserved for future use.

[00384] The 1-byte Inter-Pupillary Distance (IPD) Minimum field specifies the minimum inter-pupillary distance in millimeters (mm). If this value is zero then the minimum inter-pupillary distance is not specified. The 1-byte Inter-Pupillary Distance (IPD) Maximum field specifies the maximum inter-pupillary distance in millimeters (mm). If this value is zero then the maximum inter-pupillary distance is not specified.

[00385] The Points of Field Curvature List field contains a list of 25 2-byte parameters that specify the focal distance in thousandths of a diopter (1/m) with a range of 1 to 65535 (e.g. 1 is 0.001 diopters and 65535 is 65.535 diopters). The 25 elements in the Points of Field Curvature List are labeled A1 through E5 as shown in FIG. 82. The points are to be evenly distributed over the active area of the display. Column C corresponds to the vertical axis of the display and row 3 corresponds to the horizontal

axis of the display. Columns A and E correspond to the left and right edges of the display, respectively. And rows 1 and 5 correspond to the top and bottom edges of the display, respectively. The order of the 25 points in the list is: A1, B1, C1, D1, E1, A2, B2, C2, D2, E2, A3, B3, C3, D3, E3, A4, B4, C4, D4, E4, A5, B5, C5, D5, E5.

[00386] The CRC field contains a CRC of all bytes in the packet including the Packet Length.

38. Client Error Report Packet

[00387] The Client Error Report Packet acts as a mechanism or means for allowing a client to provide a list of operating errors to the host. The client may detect a wide range of errors in the course of its normal operation as a result of receiving certain commands from the host. Examples of these errors include: the client may have been commanded to operate in a mode that it does not support, the client may have received a packet containing certain parameters that are out of range or are beyond the capability of the client, the client may have been commanded to enter a mode in an improper sequence. The Client Error Report Packet may be used to detect errors during normal operation, but is most useful to the system designer and integrator to diagnose problems in development and integration of host and client systems. A client indicates its ability to send a Client Error Report Packet using a parameter value of 142 in the Valid Parameter Reply List of the Valid Status Reply List Packet.

[00388] The format of one embodiment of a Client Error Report Packet is shown generally in FIG. 83. As seen in FIG. 83, a Client Error Report Packet is structured to have Packet Length, Packet Type, cClient ID, Number of List Items, Error Code List, and CRC fields. A Packet Type value of 142 identifies a packet as a Client Error Report Packet. The cClient ID field is reserved for future use and is generally set to zero for now. The Number of List Items field (2 bytes) specifies the number of items in the following Error Code List. The Error Code List field (here 8 bytes) is a list containing one or more Error Report List items. The format of a single Error Report List item is shown in FIG. 84.

[00389] In one embodiment, as shown in FIG. 84, each Error Report List Item is exactly 4 bytes in length, and has a structure in one embodiment comprising: a 2-byte Display Error Code field that specifies the type of error being reported, a 2-byte Error Sub-code field specifies a greater level of detail regarding the error defined by the Client Error

Code packet. The specific definition of each Client Error Code is defined by the manufacturer of the client. An Error Sub-code does not have to be defined for every Display Error Code, and in those cases where the Error Sub-code is not defined the value is set to zero. The specific definition of each Error Sub-code is defined by the manufacturer of the client.

39. *Client Identification Packet*

[00390] The Client Identification Packet allows a client to return identifying data in response to a Request Specific Status Packet. In one embodiment, a client indicates an ability to send the Client Identification Packet using a parameter value of 144 in the Valid Parameter Reply List of the Valid Status Reply List Packet. It is useful for the host to be able to determine the client device manufacturer name and model number by reading this data from the client. The information may be used to determine if the client has special capabilities that cannot be described in the Client Capability Packet. There are potentially two methods, means, or mechanisms for reading identification information from the client. One is through use of the Client Capability Packet, which contains fields similar to those in the base EDID structure. The other method is through use of the Client Identification Packet that contains a richer set of information compared to the similar fields in the Client Capability Packet. This allows a host to identify manufacturers that have not been assigned a 3-character EISA code, and allows serial numbers to contain alphanumeric characters.

[00391] The format of one embodiment of a Client Identification Packet is shown generally in FIG. 85. As seen in FIG. 85, a Client Identification Packet is structured to have Packet Length, Packet Type, cClient ID, Week of Mfr., Year of Mfr., Length of Mfr Name, Length of Product Name, Length of Serial Number, Manufacturer Name String, Product Name String, Serial Number String, and CRC fields.

[00392] The 2 byte Packet Type field contains a value that identifies the packet as a Client Identification Packet. This value is selected to be 144 in one embodiment. The cClient ID field (2 bytes) again is reserved for future use for the Client ID, and is generally set to zero. The CRC field (2 bytes) contains a 16-bit CRC of all bytes in the packet including the Packet Length.

[00393] A 1-byte Week of Manufacture field contains a value that defines the week of manufacture of the display. In at least one embodiment, this value is in the range of 1 to

53 if it is supported by the client. If this field is not supported by the client, then it is generally set to zero. A 1-byte Year of Manufacture field contains a value that defines the year of manufacture of the client (display). This value is an offset from the year 1990 as a starting point, although other base years could be used. Years in the range of 1991 to 2245 can be expressed by this field. Example: the year 2003 corresponds to a Year of Manufacture value of 13. If this field is not supported by the client it should be set to a value of zero.

[00394] The Length of Mfr Name, Length of Product Name, and Length of Serial Number fields each contain 2-byte values that specify the length of the Manufacturer Name String field including any null termination or null pad characters, the length of the Product Name String field including any null termination or null pad characters, and the length of the Serial Number String field including any null termination or null pad characters, respectively.

[00395] The Manufacturer Name String, Product Name String, and Serial Number String fields each contain a variable number of bytes specified by the Length Mfr Name, Product Name, and Serial Number fields, respectively, that contain an ASCII string that specifies the manufacturer, product name, and alphanumeric serial number of the display, respectively. Each of these strings is terminated by at least one null character.

40. Alternate Display Capability Packet

[00396] The Alternate Display Capability Packet is used as a means, structure, or method to indicate the capability of alternate displays attached to the MDDI client controller. It is sent in response to a Request Specific Status Packet. When prompted, a client device sends an Alternate Display Capability Packet for each alternate display that is supported. If a client has more than one alternate display, then the client should send, generate, or provide multiple Alternate Display Capability Packets, one for each display, in response to a single Request Specific Status Packet, although some configurations can use multiple Request Specific Status Packets as desired, although this is less efficient. The client may send Alternate Display Capability Packets in what can be referred to as a "non-sequential order" based on the value of the Alt Display Number field. The client can indicate an ability to send the Alternate Display Capability Packet via a parameter value of 145 in the Valid Parameter Reply List of the Valid Status Reply List Packet.

[00397] For MDDI systems operated in internal mode it may be common to have more than one display connected to an MDDI client controller. An example application is a mobile phone with a large display on the inside of the flip and a smaller display on the outside. It is not necessary for an internal mode client to return an Alternate Display Capability Packet for two potential reasons. First, the host may already be programmed or otherwise informed of the capabilities during manufacture since they are used in a common device or housing. Second, due to assembly of the two, the client cannot easily be disconnected or separated from a connection to the host, and the host may contain a hard-coded copy of the client capabilities, or at least know they do not change with a change in client, as otherwise might occur.

[00398] The Number of Alt Displays field of the Client Capability Packet is used to report that more than one display is attached and the Alternate Display Capability Packet reports the capability of each alternate display. The video stream packet contains 4 bits in the Pixel Data Attributes field to address each alternate display in the client device.

[00399] The format of one embodiment of a Alternate Display Capability Packet is shown generally in FIG. 89. As seen in FIG. 86, an Alternate Display Capability Packet is structured to have Packet Length, Packet Type, cClient ID, Alt Display Number, Reserved 1, Bitmap Width, Bitmap Height, Display Window Width, Display Window Height, Color Map RGB Width, RGB Capability, Monochrome Capability, Reserved 2, Y Cb Cr Capability, Display Feature Capability, Reserved 3, and CRC fields. A Packet Type value of 145 identifies a packet as a Alternate Display Capability Packet. The cClient ID field is reserved for a Client ID for future use and generally set to zero.

[00400] The Alt Display Number field uses 1 byte to indicate the identity of the alternate display with an integer in the range of 0 to 15. The first alternate display is typically designated as number 0 and the other alternate displays are identified with unique Alt Display Number values with the largest value used being the total number of alternate displays minus 1. Values larger than the total number of alternate displays minus 1 are not used. Example: a mobile phone having a primary display and a caller-ID display connected to an MDDI client has one alternate display, so the Alt Display Number of the caller-ID display is zero and the Number of Alt Displays field of the Client Capability Packet has a value of 1.

- [00401] The Reserved 1 field (1 byte) is reserved for future use. All bits in this field are set to zero. One purpose of this field is to cause all subsequent 2 byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.
- [00402] The Bitmap Width field uses 2 bytes that specify the width of the bitmap expressed as a number of pixels. The Bitmap Height field uses 2 bytes that specify the height of the bitmap expressed as a number of pixels. The Display Window Width field uses 2 bytes that specify the width of the display window expressed as a number of pixels. The Display Window Height field uses 2 bytes that specify the height of the display window expressed as a number of pixels.
- [00403] The Color Map RGB Width field uses 2 bytes that specify the number of bits of the red, green, and blue color components that can be displayed in the color map (palette) display mode. A maximum of 8 bits for each color component (red, green, and blue) can be used. Even though 8 bits of each color component are sent in the Color Map Packet, only the number of least significant bits of each color component defined in this field are used. If the display client cannot use the color map (palette) format then this value is zero. The color map RGB Width word is composed of three separate unsigned values:
- [00404] Bits 3 through 0 define the maximum number of bits of blue in each pixel with values of 0 to 8 being considered valid. Bits 7 through 4 define the maximum number of bits of green in each pixel with values of 0 to 8 being considered valid. Bits 11 through 8 define the maximum number of bits of red in each pixel with values of 0 to 8 being considered valid. Bits 14 through 12 are reserved for future use and are generally set to zero. Bit 15 is used to indicate the ability of a client to accept Color Map pixel data in packed or unpacked format. When Bit 15 is set to a logic-one level, this indicates that the client can accept Color Map pixel data in either packed or unpacked format. If bit 15 is set to a logic-zero, this indicates that the client can accept Color Map pixel data only in unpacked format.
- [00405] RGB Capability field uses 2 bytes to specify the number of bits of resolution that can be displayed in RGB format. In one embodiment, if the client cannot use the RGB format then this value is set equal to zero. The RGB Capability word is composed of three separate unsigned values: Bits 3 through 0 define the maximum number of bits of blue (the blue intensity) in each pixel, Bits 7 through 4 define the maximum number of bits of green (the green intensity) in each pixel, and Bits 11 through 8 define the

maximum number of bits of red (the red intensity) in each pixel. Bits 14 through 12 are reserved for future use and are set to zero. Bit 15 is used to indicate the ability of a client to accept RGB pixel data in packed or unpacked format. When Bit 15 is set to a logic-one level, this indicates that the client can accept RGB pixel data in either packed or unpacked format. If bit 15 is set to a logic-zero, this indicates that the client can accept RGB pixel data only in unpacked format.

[00406] The 1 byte Monochrome Capability field contains a value or information to specify the number of bits of resolution that can be displayed in monochrome format. If the client cannot use the monochrome format then this value is set equal to zero. Bits 6 through 4 are reserved for future use and are generally set to zero. Bits 3 through 0 define the maximum number of bits of grayscale that can exist in each pixel. These four bits make it possible to specify that each pixel consists of 1 to 15 bits. If the value is zero then the monochrome format is not supported by the client. Bit 7 when set to one indicates that the client can accept monochrome pixel data in either packed or unpacked format. If bit 7 is set to zero this indicates that the client can accept monochrome pixel data only in unpacked format.

[00407] The Reserved 2 field is a 1 byte wide field reserved for future use and generally has all bits in this field set to logic-zero level. In one embodiment, one purpose of this field is to cause all subsequent 2 byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.

[00408] A 2-byte Y Cb Cr Capability field specifies the number of bits of resolution that can be displayed in Y Cb Cr format. If the client cannot use the Y Cb Cr format then this value is zero. The Y Cb Cr Capability word is composed of three separate unsigned values: Bits 3 through 0 define the maximum number of bits that specify the Cb sample, Bits 7 through 4 define the maximum number of bits that specify the Cr sample, Bits 11 through 8 define the maximum number of bits that specify the Y sample, and Bits 14 through 12 are reserved for future use and are set to zero. Bit 15 when set to one indicates that the client can accept Y Cb Cr pixel data in either packed or unpacked format. If bit 15 is set to zero this indicates that the client can accept Y Cb Cr pixel data only in unpacked format.

[00409] A 2 byte Bayer Capability field specifies the number of bits of resolution, pixel group, and pixel order that can be transferred in Bayer format. If the client cannot use the Bayer format then this value is set at zero. The Bayer Capability field is composed

of the following values: Bits 3 through 0 define the maximum number of bits of intensity that exist in each pixel, Bits 5 through 4 define the pixel group pattern that may be required. Bits 8 through 6 define a pixel order that is required, and Bits 14 through 9 are reserved for future use and are set to zero. Bit 15 when set to one indicates that the client can accept Bayer pixel data in either packed or unpacked format. If bit 15 is set to zero, this indicates that the client can accept Bayer pixel data only in unpacked format.

- [00410] The 2-byte CRC field contains a 16-bit CRC of all bytes in the packet including the Packet Length.

41. Register Access Packet

- [00411] The Register Access Packet provides either a host or a client with a means, mechanism, or method to access configuration and status registers in the opposite end of the MDDI link. The registers are likely to be unique for each display or device controller. These registers already exist in many displays that require setting configurations, modes of operation, and have other useful and necessary settings. The Register Access Packet allows the MDDI host or client to both write to a register and request to read a register using the MDDI link. When the host or client requests to read a register the opposite end should respond by sending the register data in the same packet type, but also by indicating that this is the data read from a particular register with the use of the Read/Write Info field. The Register Access Packet may be used to read or write multiple registers by specifying a register count greater than 1. A client indicates an ability to support the Register Access Packet using bit 22 of Client Feature Capability field of the Client Capability Packet. The client will use the encapsulation packet to send the Register Access Packet, therefore presenting what appears as a packet within a packet configuration or structure.

- [00412] The format of one embodiment of a Register Access Packet is shown generally in FIG. 87. As seen in FIG. 87, a Register Access Packet is structured to have Packet Length, Packet Type, bClient ID, Read/Write Flags, Register Address, Parameter CRC, Register Data List and Register Data CRC fields. A Packet Type value of 146 identifies a packet as Register Access Packet. The bClient ID field is reserved for future use and is generally set to zero for now.

[00413] The 2-byte Read/Write Flags field specifies the specific packet as either a write, or a read, or a response to a read, and provides a count of the data values.

[00414] Bits 15 through 14 act as Read/Write Flags. If Bits[15:14] are '00' then this packet contains data to be written to a register addressed by the Register Address field. The data to be written to the specified registers is contained in the Register Data List field. If Bits[15:14] are '10' then this is a request for data from one or more registers addressed by the Register Address field. If Bits[15:14] are '11' then that packet contains data that was requested in response to a Register Access Packet having bits 15:14 of the Read/Write Flags set to '10'. The Register Address field contains the address of the register corresponding to the first Register Data List item, and the Register Data List field contains data that was read from the address or addresses. If Bits[15:14] are '01' this is treated as an invalid value, this value is reserved for future use and is not used at this time, but those skilled in the art will understand how to employ it for future applications.

[00415] Bits 13:0 use a 14-bit unsigned integer to specify the number of 32-bit Register Data items to be transferred in the Register Data List field. If bits 15:14 equal '00' then bits 13:0 specify the number of 32-bit register data items that are contained in the Register Data List field to be written to registers starting at the register specified by the Register Address field. If bits 15:14 equal '10' then bits 13:0 specify the number of 32-bit register data items that the receiving device sends to a device requesting that the registers be read. The Register Data List field in this packet contains no items and is of zero length. If bits 15:14 equal '11' then bits 13:0 specify the number of 32-bit register data items that have been read from registers that are contained in the Register Data List field. Bits 15:14 are not currently set equal to '01', which is considered an invalid value, and otherwise reserved for future designations or use.

[00416] The Register Address field uses 4 bytes to indicate the register address that is to be written to or read from. For addressing registers whose addressing is less than 32 bits, the upper bits are set to zero.

[00417] The 2-byte Parameter CRC field contains a CRC of all bytes from the Packet Length to the Register Address. If this CRC fails to check then the entire packet is discarded.

[00418] The Register Data List field contains a list of 4-byte register data values to be written to client registers or values that were read from client device registers.

- [00419] The 2-byte Register Data CRC field contains a CRC of only the Register Data List. If this CRC fails to check then the Register Data may still be used, but the CRC error count is incremented.

D. Packet CRC

- [00420] The CRC fields appear at the end of the packets and sometimes after certain more critical parameters in packets that may have a significantly large data field, and thus, an increased likelihood of errors during transfer. In packets that have two CRC fields, the CRC generator, when only one is used, is re-initialized after the first CRC so that the CRC computations following a long data field are not affected by the parameters at the beginning of the packet.
- [00421] There is a remote possibility for packets containing multiple bit errors to produce a good CRC. The probability of detecting a good CRC on a packet with errors approaches 7.6×10^{-6} on very long packets containing many errors. By design, the MDDI link will have a very low or zero error rate. The CRC is intended to be used to monitor the health of the link, and is not intended to detect errors on specific packets to determine whether packets should be retransmitted.
- [00422] In an exemplary embodiment, the polynomial used for the CRC calculation is known as the CRC-16, or $X^{16} + X^{15} + X^2 + X^0$. A sample implementation of a CRC generator and checker 3600 useful for implementing the invention is shown in FIG. 36. In FIG. 36, a CRC register 3602 is initialized to a value of 0x0001 just prior to transfer of the first bit of a packet which is input on the Tx_MDDI_Data_Before_CRC line, then the bytes of the packet are shifted into the register starting with the LSB first. Note that the register bit numbers in this figure correspond to the order of the polynomial being used, and not the bit positions used by the MDDI. It is more efficient to shift the CRC register in a single direction, and this results in having CRC bit 15 appear in bit position 0 of the MDDI CRC field, and CRC register bit 14 in MDDI CRC field bit position 1, and so forth until MDDI bit position 14 is reached.
- [00423] As an example, if the packet contents for the Client Request and Status Packets are: 0x000c, 0x0046, 0x000, 0x0400, 0x00, 0x00, 0x0000 (or represented as a sequence of bytes as: 0x0c, 0x00, 0x46, 0x00, 0x00, 0x00, 0x00, 0x04, 0x00, 0x00, 0x00, 0x00), and are submitted using the inputs of the multiplexors 3604 and 3606, and AND gate

3608, the resulting CRC output on the Tx_MDDI_Data_With_CRC line is 0xd9aa (or represented as a sequence as 0xaa, 0xd9).

[00424] When CRC generator and checker 3600 is configured as a CRC checker, the CRC that is received on the Rx_MDDI_Data line is input to multiplexor 3604 and exclusive-OR (XOR) gate 3612, and is compared bit by bit with the value found in the CRC register using NOR gate 3610, AND gate 3608, and AND gate 3614. If there are any errors, as output by AND gate 3614, the CRC is incremented once for every packet that contains a CRC error by connecting the output of gate 3614 to the input of register 3602. Note that the example circuit shown in the diagram of FIG. 36 can output more than one CRC error signal within a given CHECK_CRC_NOW window (see FIG. 37B). Therefore, the CRC error counter generally only counts the first CRC error instance within each interval where CHECK_CRC_NOW is active. If configured as a CRC generator the CRC is clocked out of the CRC register at the time coinciding with the end of the packet.

[00425] The timing for the input and output signals, and the enabling signals, is illustrated graphically in FIGs. 37A and 37B. The generation of a CRC and transmission of a packet of data are shown in FIG. 37A with the state (0 or 1) of the Gen_Reset, Check_CRC_Now, Generate_CRC_Now, and Sending_MDDI_Data signals, along with the Tx_MDDI_Data_Before_CRC and Tx_MDDI_Data_With_CRC signals. The reception of a packet of data and checking of the CRC value are shown in FIG. 37B, with the state of the Gen_Reset, Check_CRC_Now, Generate_CRC_Now, and Sending_MDDI_Data signals, along with the Rx_MDDI_Data and CRC error signals.

E. Error Code Overload for Packet CRC

[00426] Whenever only data packets and CRC are being transferred between the host and client, there are no error codes being accommodated. The only error is a loss of synchronization. Otherwise, one has to wait for the link to timeout from a lack of a good data transfer path or pipeline and then reset the link and proceed. Unfortunately, this is time consuming and somewhat inefficient.

[00427] For use in one embodiment, a new technique has been developed in which the CRC portion of packets is used to transfer error code information. This is generally shown in FIG. 65. That is, one or more error codes are generated by the processors or

devices handling the data transfer which indicate specific predefined errors or flaws that might occur within the communication processing or link. When an error is encountered, that the appropriate error code is generated and transferred using the bits for the CRC of a packet. That is, the CRC value is overloaded, or overwritten, with the desired error code, which can be detected on the receiving end by an error monitor or checker that monitors the values of the CRC field. For those cases in which the error code matches the CRC value for some reason, the compliment of the error is transferred to prevent confusion.

[00428] In one embodiment, to provide a robust error warning and detection system, the error code may be transferred several times, using a series of packets, generally all, that are transferred or sent after the error has been detected. This occurs until the point at which the condition creating the error is cleared from the system, at which point the regular CRC bits are transferred without overloading by another value.

[00429] This technique of overloading the CRC value provides a much quicker response to system errors while using a minimal amount of extra bits or fields.

[00430] As shown in FIG. 66, a CRC overwriting mechanism or apparatus 6600 is shown using an error detector or detections means 6602, which can form part of other circuitry previously described or known, detects the presence or existence of errors within the communication link or process. An error code generator or means 6604, which can be formed as part of other circuitry or use techniques such as look up tables to store pre-selected error messages, generates one or more error codes to indicate specific predefined errors or flaws that have been detected as occurring. It is readily understood that devices 6602 and 6604 can be formed as a single circuit or device as desired, or as part of a programmed sequence of steps for other known processors and elements.

[00431] A CRC value comparator or comparison means 6606 is shown for checking to see if the selected error code or codes are the same as the CRC value being transferred. If that is the case then a code compliment generator or generation means or device is used to provide the compliment of the error codes as to not be mistaken as the original CRC pattern or value and confuse or complicate the detection scheme. An error code selector or selection means element or device 6610 then selects the error code or value it is desired to insert or overwrite, or their respective compliments as appropriate. An error code CRC over-writer or over writing mechanism or means 6612 is a device that

receives the data stream, packets, and the desired codes to be inserted and overwrites the corresponding or appropriate CRC values, in order to transfer the desired error codes to a receiving device.

[00432] As mentioned, the error code may be transferred several times, using a series of packets, so the over-writer 6612 may utilize memory storage elements in order to maintain copies of the codes during processing or recall these codes from previous elements or other known storage locations which can be used to store or hold their values as needed, or as desired.

[00433] The general processing the overwriting mechanism of FIG. 66 is implementing is shown in additional detail in FIGs. 67A and 67B. In 67A, an error, one or more, is detected in step 6702 in the communication data or process and an error code is selected in step 6704 to indicate this condition. At the same time, or at an appropriate point, the CRC value to be replaced is checked in a step 6706, and compared to the desired error code in step 6708. The result of this comparison, as discussed earlier, is a determination as to whether or not the desired code, or other representative values, will be the same as the CRC value present. If this is the case, then processing proceeds to a step 6712 where the compliment, or in some cases another representative value, as desired, is selected as the code to insert. Once it has been determined what error codes or values are to be inserted in steps 6710 and 6714, that appropriate code is selected for insertion. These steps are illustrated as separate for purposes of clarity but generally represent a single choice based on the output of the step 6708 decision. Finally, in step 6716 the appropriate values are overwritten in the CRC location for transfer with the packets being targeted by the process.

[00434] On the packet reception side, as shown in FIG. 67B, the packet CRC values are being monitored in a step 6722. Generally, the CRC values are being monitored by one or more processes within the system to determine if an error in data transfer has occurred and whether or not to request a retransmission of the packet or packets, or to inhibit further operations and so forth, some of which is discussed above. As part of such monitoring the information can also be used to compare values to known or pre-selected error codes, or representative values and detect the presence of errors. Alternatively, a separate error detection process and monitor can be implemented. If a code appears to be present it is extracted or otherwise noted in step 6724 for further processing. A determination can be made in step 6726 as to whether or not this is the

actual code or a compliment, in which case an additional step 6728 is used to translate the value to the desired code value. In either case the resulting extracted code, compliment, or other recovered values are then used to detect what error has occurred from the transmitted code in step 6730.

V. Link Hibernation

[00435] The MDDI link can enter the hibernation state quickly and wake up from hibernation quickly. This responsiveness allows a communicating system or device to force the MDDI link into hibernation frequently to reduce power consumption, since it can wake up again for use very quickly. In one embodiment, as an external mode client wakes up from hibernation for the first time it does so at a data rate and with strobe pulse timing that is consistent with a 1 Mbps rate, that is, the MDDI_Stb pair should toggle at a 500 kHz rate. Once characteristics of the client have been discovered by or communicated to the host, then the host may wake up the link at generally any rate from 1 Mbps to the maximum rate at which the client can operate. Internal mode clients may wake up at any rate at which both the host and client can operate. This is also generally applicable to the first time an internal mode client wakes up.

[00436] In one embodiment, when the link wakes up from hibernation the host and client exchange a sequence of pulses. These pulses can be detected using low-speed line receivers that consume only a fraction of the current as the differential receivers required to receive the signals at the maximum link operating speed. Either the host or client can wake up the link, so the wake-up protocol is designed to handle possible contention that can occur if both host and client attempt to wake up simultaneously.

[00437] During the hibernation state the MDDI_Data and MDDI_Stb differential drivers are disabled in the high-impedance state and the differential voltage across all differential pairs is zero volts. The differential line receivers used to detect the sequence of pulses during the wake-up from hibernation have an intentional voltage offset. In one embodiment, the threshold between a logic-one and logic-zero level in these receivers is approximately 125 mV. This causes an un-driven differential pair to be seen as a logic-zero level during the link wake-up sequence.

[00438] In order to enter a Hibernation State, the host sends 64 MDDI_Stb cycles after the CRC of the Link Shutdown Packet. The host disables the MDDI_Data0 output of

the host in the range of 16 to 56 MDDI_Stb cycles (including output disable propagation delays) after the CRC. The host finishes sending the 64 MDDI_Stb cycles after the CRC of the Link Shutdown packet before it initiates the wake-up sequence. In one embodiment, the host-initiated wake-up is defined as the host having to wait at least 100 nsec after MDDI_Data0 reaches a valid logic-one level before driving pulses on MDDI_Stb. In one embodiment, the client waits at least 60 MDDI_Stb cycles after the CRC of the Link Shutdown Packet before it drives MDDI_Data0 to a logic-one level to attempt to wake-up the host.

[00439] In order to "wake-up" from a Hibernation State, several actions or processes are undertaken. When the client, here a display, needs data or communication, service, from the host it generates a request pulse by driving the MDDI_Data0 line to a logic-one state for around 70 to 1000 μ sec, while MDDI_Stb is inactive and keeps MDDI_Data0 driven to a logic-one level for about 70 MDDI_Stb cycles (over a range of 60 to 80) after MDDI_Stb becomes active, although other periods can be used as desired. The client then disables the MDDI_Data0 driver by placing it into a high-impedance state.

[00440] If MDDI_Stb is active during hibernation, although unlikely, then the client might only drive MDDI_Data0 to a logic-one state for about 70 MDDI_Stb cycles (over a range of 60 to 80). This action causes the host to start or restart data traffic on the forward link (208) and to poll the client for its status.

[00441] The host must detect the presence of the request pulse and begins the startup sequence of first driving the MDDI_Stb to logic-zero level and MDDI_Data0 to a logic-high level for at least around 200 nsec. And then while toggling MDDI_Stb continue to drive MDDI_Data0 to a logic-one level for about 150 MDDI_Stb cycles (a range of 140 to 160) and to logic-zero for about 50 MDDI_Stb cycles. The client should not send a service request pulse if it detects MDDI_Data0 in the logic-one state for more than 80 MDDI_Stb cycles. When the client has detected MDDI_Data0 at a logic-one level for 60 to 80 MDDI_Stb cycles it begins to search for the interval where the host drives MDDI_Data0 to a logic-zero level for 50 MDDI_Stb cycles. After the host drives MDDI_Data0 to a logic-zero level for a duration of 50 MDDI_Stb cycles, then the host starts sending packets on the link. The first packet sent is a Sub-frame Header Packet. The client begins to look for the Sub-frame Header Packet after MDDI_Data0 is at a logic-zero level for 40 MDDI_Stb cycles of the 50 cycle interval. The nature of

selection of the times and tolerances of time intervals related to the hibernation processing and start up sequence are discussed further below. (See FIGs. 68A-C below.)

[00442] The host may initiate the wake-up by first enabling MDDI_Stb and simultaneously drive it to a logic-zero level. MDDI_Stb should not be driven to a logic-one level until pulses are output as described below. After MDDI_Stb reaches a logic-zero level the host enables MDDI_Data0 and simultaneously drives it to a logic-one level. MDDI_Data0 should not be driven to a logic-zero level during the wake-up process until the interval where it is driven to a logic-zero level for an interval of 50 MDDI_Stb pulses as described below. The host should wait at least 200 nsec after MDDI_Data0 reaches a valid logic-one level before driving pulses on MDDI_Stb. This timing relationship occurs while considering the worst-case output enable delays. This substantially guarantees that a client has sufficient time to fully enable its MDDI_Stb receiver after being awakened by a logic-one level on MDDI_Data0 that was driven by the host.

[00443] An example of the processing steps for a typical client service request event 3800 with no contention is illustrated in FIG. 38, where the events are labeled for convenience in illustration using the letters A, B, C, D, E, F, and G. The process commences at point A when the host sends a Link Shutdown Packet to the client device to inform it that the link will transition to a low-power hibernation state. In a next step, the host enters the low-power hibernation state by disabling the MDDI_Data0 driver and setting the MDDI_Stb driver to a logic zero, as shown at point B. MDDI_Data0 is driven to a logic-zero level by a high-impedance bias network. After some period of time, the client sends a service request pulse to the host by driving MDDI_Data0 to a logic one level as seen at point C. The host still asserts the logic-zero level using the high-impedance bias network, but the driver in the client forces the line to a logic one level. Within 50 μ sec, the host recognizes the service request pulse, and asserts a logic one level on MDDI_Data0 by enabling its driver, as seen at point D. The client then ceases from attempting to assert the service request pulse, and the client places its driver into a high-impedance state, as seen at point E. The host drives MDDI_Data0 to a logic-zero level for 50 μ sec, as shown at point F, and also begins to generate MDDI_Stb in a manner consistent with the logic-zero level on MDDI_Data0. The client begins to look for the Sub-frame Header Packet after MDDI_Data0 is at a logic-zero level for 40

MDDI_Stb cycles. After asserting MDDI_Data0 to a logic-zero level and driving MDDI_Stb for 50 μ sec, the host begins to transmit data on the forward link by sending a Sub-frame Header Packet, as shown at point G.

[00444] A similar example is illustrated in FIG. 39 where a service request is asserted after the link restart sequence has begun, and the events are again labeled using the letters A, B, C, D, E, F, and G. This represents a worst case scenario where a request pulse or signal from the client comes closest to corrupting the Sub-frame Header Packet. The process commences at point A when the host again sends a Link Shutdown Packet to the client device to inform it that the link will transition to a low-power hibernation state. In a next step, the host enters the low-power hibernation state by disabling the MDDI_Data0 driver and setting the MDDI_Stb driver to a logic-zero level, as shown at point B. As before, MDDI_Data0 is driven to a logic-zero level by a high-impedance bias network. After a period of time, the host begins the link restart sequence by driving MDDI_Data0 to a logic-one level for 150 μ sec as seen at point C. Prior to 50 μ sec passing after the link restart sequence begins the display also asserts MDDI_Data0 for a duration of 70 μ sec, as seen at point D. This happens because the display has a need to request service from the host and does not recognize that the host has already begun the link restart sequence. The client then ceases attempting to assert the service request pulse, and the client places its driver into a high-impedance state, as seen at point E. The host continues to drive MDDI_Data0 to a logic-one level. The host drives MDDI_Data0 to a logic-zero level for 50 μ sec, as shown at point F, and also begins to generate MDDI_Stb in a manner consistent with the logic zero level on MDDI_Data0. After asserting MDDI_Data0 to a logic-zero level, and driving MDDI_Stb for 50 μ sec, the host begins to transmit data on the forward link by sending a Sub-frame Header Packet, as shown at point G.

[00445] From the above discussion, one sees that the prior solution involved having the host go through two states as part of a wakeup sequence. For the first state, the host drives the MDDI_Data0 signal high for 150 μ s, and then drives the MDDI_Data0 signal low for 50 μ s while activating the MDDI_Stb line, and then begins to transmit MDDI packets. This process works well to advance the state of the art in terms of data rates achievable using the MDDI apparatus and methods. However, as stated earlier, more speed in terms of reduced response time to conditions or being able to more quickly

select the next step or process, are the ability to simplify processing or elements, are always in demand.

[00446] Applicants have discovered a new inventive approach to wake-up processing and timing in which the host uses a clock cycle based timing for the signal toggling. In this configuration, the host starts toggling MDDI_Stb from 0 to 10 μ sec after the host drives the MDDI_Data0 signal high at the beginning of the wake-up sequence, and does not wait until the signal is driven low. During a wake-up sequence, the host toggles MDDI_Stb as though the MDDI_Data0 signal were always at a logic-zero level. This effectively removes the concept of time from the client side, and the host changes from the prior 150 μ s and 50 μ s periods for the first two states, to 150 clock cycles and 50 clock cycles, for these periods.

[00447] The host now becomes responsible for driving that data line high, and within 10 clock cycles starting to transmit a strobe signal as if the data line was zero. After the host has driven the data line high for 150 clock cycles, the host drives the data line low for 50 clock cycles while continuing to transmit the strobe signal. After it has completed both of these processes, the host can begin to transmit the first sub-frame header packet.

[00448] On the client side, the client implementation can now use the generated clock to calculate the number of clock cycles that the data line is first high, and then low. The number of clock cycles that need to occur in both the data line driven high state is 150 and data line driven low state is 50. This means that for a proper wakeup sequence, the client should be able to count at least 150 continuous clock cycles of the data line being high, followed by at least 50 continuous clock cycles of the data line being low. Once these two conditions are met, the client can begin to search for the unique word of the first sub-frame. A break in this pattern is used as a basis to return the counters to an initial state in which the client again looks for the first 150 continuous clock cycles of the data line being high.

[00449] A client implementation of the invention for host based wakeup from hibernation is very similar to the initial start-up case except that the clock rate is not forced to start at 1Mbps, as discussed earlier. Instead the clock rate can be set to resume at whatever previous rate was active when the communication link went into hibernation. If the host begins transmission of a strobe signal as described above, the client should be able to again count at least 150 continuous clock cycles of the data line

being high, followed by at least 50 continuous clock cycles of the data line being low. Once these two conditions have been met, the client can begin the search for the unique word.

[00450] A client implementation of the invention for client based wakeup from hibernation is similar to the host based wakeup except that it starts by having the client driving the data line. The client can asynchronously drive the data line without a clock to wake up the host device. Once the host recognizes that the data line is being driven high by the client, it can begin its wakeup sequence. The client can count the number of clock cycles generated by the host starting or during its wakeup process. Once the client counts 70 continuous clock cycles of the data being high, it can stop driving the data line high. At this point, the host should already be driving the data line high as well. The client can then count another 80 continuous clock cycles of the data line being high to reach the 150 clock cycles of the data line being high, and can then look for 50 clock cycles of the data line being low. Once these three conditions have been met the client can begin to look for the unique word.

[00451] An advantage of this new implementation of wake-up processing is that it removes the need for a time measuring device. Whether this is an oscillator, or capacitor discharge circuit, or other such known devices, the client no longer needs such external devices to determine the start up conditions. This saves money and circuit area when implementing controllers, counters, and so forth on a client device board. While this may not be as advantageous to the client, for the host, this technique should also potentially simplify the host in terms of Very High Density Logic (VHDL) being used for core circuitry. The power consumption of using the data and strobe lines as the wakeup notification and measurement source will also be lower since no external circuitry will need to be running for the core elements to be waiting for a host based wakeup. The number of cycles or clock periods used is exemplary and other periods can be used as will be apparent to one skilled in the art.

[00452] An advantage of this new implementation of wake-up processing is that it removes the need for a time measuring device. Whether this is an oscillator, or capacitor discharge circuit, or other such known devices, the client no longer needs such external devices to determine the start up conditions. This saves money and circuit area when implementing controllers, counters, and so forth on a client device board. While this may not be as advantageous to the client, for the host, this technique should also

potentially simplify the host in terms of VHDL being used for core circuitry. The power consumption of using the data and strobe lines as the wakeup notification and measurement source will also be lower since no external circuitry will need to be running for the core elements to be waiting for a host based wakeup.

[00453] To clarify and illustrate the operation of this new technique, the timing of MDDI_Data0, MDDI_Stb, and various operations relative to the clock cycles are shown in FIGs. 68A, 68B, and 68C.

[00454] An example of the processing steps for a typical Host-initiated Wake-up with no contention is illustrated in FIG. 68A, where the events are again labeled for convenience in illustration using the letters A, B, C, D, E, F, and G. The process commences at point A when the host sends a Link Shutdown Packet to the client device to inform it that the link will transition to a low-power hibernation state. In a next step, point B, the host toggles MDDI_Stb for about 64 cycles (or as desired for system design) to allow processing by the client to be completed prior to stopping MDDI_Stb from toggling, which stops the recovered clock in the client device. The host also initially sets MDDI_Data0 to logic-zero level and then disables the MDDI_Data0 output in the range of 16 to 48 cycles (generally including output disable propagation delays) after the CRC. It may be desirable to place high-speed receivers for MDDI_Data0 and MDDI_Stb in the client in a low power state some time after the 48 cycles after the CRC and prior to the next stage (C). The client places its high-speed receivers for MDDI_Data0 and MDDI_Stb into hibernation any time after the rising edge of the 48th MDDI_Stb cycle after the CRC of the Link Shutdown Packet. It is recommended that the client place its high-speed receivers for MDDI_Data0 and MDDI_Stb into hibernation before the rising edge of the 64th MDDI_Stb cycle after the CRC of the Link Shutdown Packet.

[00455] The host enters the low-power hibernation state at point or step C, by disabling the MDDI_Data0 and MDDI_Stb drivers and placing a host controller in a low power hibernation state. One can also set the MDDI_Stb driver to a logic-zero level (using a high-impedance bias network) or to continue toggling during hibernation, as desired. The client is also in a low power level hibernation state.

[00456] After some period of time, the host commences the link restart sequence at point D, by enabling the MDDI_Data0 and MDDI_Stb driver output. The host drives MDDI_Data0 to a logic-one level and MDDI_Stb to a logic-zero level for as long as it

should take for the drivers to fully enable their respective outputs. The host typically waits around 200 nanoseconds after these outputs reach desired logic levels before driving pulses on MMDI_Stb. This allows the client time to prepare to receive.

[00457] With the host drivers enabled and MMDI_Data0 being driven to a logic-one level, the host begins to toggle MMDI_Stb for a duration of 150 MMDI_Stb cycles, as seen at point E. The host drives MMDI_Data0 to a logic zero level for 50 cycles, as shown at point F, and the client begins to look for the Sub-frame Header Packet after MMDI_Data0 is at a logic-zero level for 40 MMDI_Stb cycles. The host begins to transmit data on the forward link by sending a Sub-frame Header Packet, as shown at point G.

[00458] An example of the processing steps for a typical Client-initiated Wake-up with no contention is illustrated in FIG. 68B, where the events are again labeled for convenience in illustration using the letters A, B, C, D, E, F, G, H, and I. As before, the process commences at point A when the host sends a Link Shutdown Packet to inform the client that the link will transition to the low power state.

[00459] At point B, the host toggles MMDI_Stb for about 64 cycles (or as desired for system design) to allow processing by the client to be completed prior to stopping MMDI_Stb from toggling, which stops the recovered clock in the client device. The host also initially sets MMDI_Data0 to a logic-zero level and then disables the MMDI_Data0 output in the range of 16 to 48 cycles (generally including output disable propagation delays) after the CRC. It may be desirable to place high-speed receivers for MMDI_Data0 and MMDI_Stb in the client in a low power state some time after the 48 cycles after the CRC and prior to the next stage (C).

[00460] The host enters the low-power hibernation state at point or step C, by disabling the MMDI_Data0 and MMDI_Stb drivers and placing a host controller in a low power hibernation state. One can also set the MMDI_Stb driver to a logic-zero level (using a high-impedance bias network) or to continue toggling during hibernation, as desired. The client is also in a low power level hibernation state.

[00461] After some period of time, the client commences the link restart sequence at point D, by enabling the MMDI_Stb receiver, and also enabling an offset in the MMDI_Stb receiver to guarantee the state of the received version of MMDI_Stb is a logic-zero level in the client before the host enables its MMDI_Stb driver. It may be desirable for the client to enable the offset slightly ahead of enabling the receiver to

ensure the reception of a valid differential signal and inhibit erroneous signals, as desired. The Client enables the MDDI_Data0 driver while driving the MDDI_Data0 line to a logic one level. It is allowed for MDDI_Data0 and MDDI_Stb to be enabled simultaneously if the time to enable the offset and enable the standard MDDI_Stb differential receiver is less than 200 nsec.

[00462] Within about 1 msec., point E, the host recognizes the service request pulse from the client, and the host begins the link restart sequence by enabling the MDDI_Data0 and MDDI_Stb driver outputs. The host drives MDDI_Data0 to a logic-one level and MDDI_Stb to a logic-zero level for as long as it should take for the drivers to enable their respective outputs. The host typically waits around 200 nanoseconds after these outputs reach desired logic levels before driving pulses on MDDI_Stb. This allows the client time to prepare to receive.

[00463] With the host drivers enabled and MDDI_Data0 being driven to a logic-one level, the host begins outputting pulses on MDDI_Stb for a duration of 150 MDDI_Stb cycles, as seen at point F. When the client recognizes the first pulse on MDDI_Stb it disables the offset in its MDDI_Stb receiver. The client continues to drive MDDI_Data0 to a logic-one level for 70 MDDI_Stb cycles, and disables its MDDI_Data0 driver at point G. The host continues to drive MDDI_Data0 to a logic-one level for a duration of 80 additional MDDI_Stb pulses, and at point H drives MDDI_Data0 to a logic-zero level.

[00464] As seen at points G and H, the host drives MDDI_Data0 to a logic-zero level for 50 cycles, and the client begins to look for the Sub-frame Header Packet after MDDI_Data0 is at a logic-zero level for 40 MDDI_Stb cycles. After driving MDDI_Stb for a duration of 50 cycles, the host begins to transmit data on the forward link by sending a Sub-frame Header Packet, as shown at point I.

[00465] An example of the processing steps for a typical Host-initiated Wake-up with contention from the client, that is the client also wants to wake up the link, is illustrated in FIG. 68C. The events are again labeled for convenience in illustration using the letters A, B, C, D, E, F, G, H, and I. As before, the process commences at point A when the host sends a Link Shutdown Packet to inform the client that the link will transition to the low power state, proceeds to point B where MDDI_Stb is toggled for about 64 cycles (or as desired for system design) to allow processing by the client to be completed, and then to point C, where the host enters the low-power hibernation state,

by disabling the MDDI_Data0 and MDDI_Stb drivers and placing a host controller in a low power hibernation state. After some period of time, the host commences the link restart sequence at point D, by enabling the MDDI_Data0 and MDDI_Stb driver output, and begins to toggle MDDI_Stb for a duration of 150 MDDI_Stb cycles, as seen at point E.

[00466] At up to 70 MDDI_Stb cycles after point E, here point F, the client has not yet recognized that the host is driving MDDI_Data0 to a logic-one level so the client also drives MDDI_Data0 to a logic-one level. This occurs here because the client has a desire to request service but does not recognize that the host it is trying to communicate with has already begun the link restart sequence. At point G, the client ceases to drive MDDI_Data0, and places its driver into a high impedance state by disabling its output. The host continues to drive MDDI_Data0 to a logic-one level for 80 additional cycles.

[00467] The host drives MDDI_Data0 to a logic zero level for 50 cycles, as shown at point H, and the client begins to look for the Sub-frame Header Packet after MDDI_Data0 is at a logic-zero level for 40 MDDI_Stb cycles. The host begins to transmit data on the forward link by sending a Sub-frame Header Packet, as shown at point I.

VI. Interface Electrical Specifications

[00468] In the example embodiments, Data in a Non-Return-to-Zero (NRZ) format is encoded using a data-strobe signal or DATA-STB format, which allows clock information to be embedded in the data and strobe signals. The clock can be recovered without complex phase lock loop circuitry. Data is carried over a bi-directional differential link, generally implemented using a wire-line cable, although other conductors, printed wires, or transfer elements can be used, as stated earlier. The Strobe signal (STB) is carried over a uni-directional link which is driven only by the host. The strobe signal toggles value (0 or 1) whenever there is a back-to-back state, 0 or 1, that remains the same on the Data line or signal.

[00469] An example of how a data sequence such as bits "1110001011" can be transmitted using DATA-STB encoding is shown in graphical form in FIG. 40. In FIG. 40, a DATA signal 4002 is shown on the top line of a signal timing chart and a STB signal 4004 is shown on a second line, each time aligned as appropriate (common

starting point). As time passes, when there is a change of state occurring on the DATA line 4002 (signal), then the STB line 4004 (signal) maintains a previous state, thus, the first '1' state of the DATA signal correlates with the first '0' state for the STB signal, its starting value. However, if or when the state, level, of the DATA signal does not change then the STB signal toggles to the opposite state or '1' in the present example, as is the case in FIG. 40 where the DATA is providing another '1' value. That is, there is one and only one transition per bit cycle between DATA and STB. Therefore, the STB signal transitions again, this time to '0' as the DATA signal stays at '1' and holds this level or value as the DATA signal changes level to '0'. When the DATA signal stays at '1', the STB signal toggles to the opposite state or '1' in the present example, and so forth, as the DATA signal changes or holds levels or values.

[00470] Upon receiving these signals, an exclusive-OR (XOR) operation is performed on the DATA and STB signals to produce a clock signal 4006, which is shown on the bottom of the timing chart for relative comparison with the desired data and strobe signals. An example of circuitry useful for generating the DATA and STB outputs or signals from input data at the host, and then recovering or recapturing the data from the DATA and STB signals at the client, is shown in FIG. 41.

[00471] In FIG. 41, a transmission portion 4100 is used to generate and transmit the original DATA and STB signals over an intermediary signal path 4102, while a reception portion 4120 is used to receive the signals and recover the data. As shown in FIG. 41, in order to transfer data from a host to a client, the DATA signal is input to two D-type flip-flop circuit elements 4104 and 4106 along with a clock signal for triggering the circuits. The two flip-flop circuit outputs (Q) are then split into a differential pair of signals MDDI_Data0+, MDDI_Data0- and MDDI_Stb+, MDDI_Stb-, respectively, using two differential line drivers 4108 and 4110 (voltage mode). A three-input exclusive-NOR (XNOR) gate, circuit, or logic element 4112 is connected to receive the DATA and outputs of both flip-flops, and generates an output that provides the data input for the second flip-flop, which in turn generates the MDDI_Stb+, MDDI_Stb- signals. For convenience, the XNOR gate has the inversion bubble placed to indicate that it is effectively inverting the Q output of the flip-flop that generates the Strobe.

[00472] In reception portion 4120 of FIG. 41, the MDDI_Data0+, MDDI_Data0- and MDDI_Stb+, MDDI_Stb- signals are received by each of two differential line receivers 4122 and 4124, which generate single outputs from the differential signals. The outputs

of the amplifiers are then input to each of the inputs of a two-input XOR gate, circuit, or logic element 4126 which produces the clock signal. The clock signal is used to trigger each of two D-type flip-flop circuits 4128 and 4130 which receive a delayed version of the DATA signal, through delay element 4132, one of which (4128) generates data '0' values and the other (4130) data '1' values. The clock has an independent output from the XOR logic as well. Since the clock information is distributed between the DATA and STB lines, neither signal transitions between states faster than half of the clock rate. Since the clock is reproduced using the exclusive-OR processing of the DATA and STB signals, the system effectively tolerates twice the amount of skew between the input data and clock compared to the situation when a clock signal is sent directly over a single dedicated data line.

[00473] The MDDI Data pairs, MDDI_Stb+, and MDDI_Stb- signals are operated in a differential mode to maximize immunity from the negative affects of noise. Each differential pair is parallel-terminated with the characteristic impedance of the cable or conductor being used to transfer signals. Generally, all parallel-terminations reside in the client device. This is near the differential receiver for forward traffic (data sent from the host to the client), but it is at the driving end of the cable or other conductors or transfer elements for reverse traffic (data sent from the client to the host). For reverse traffic, the signal is driven by the client, reflected by the high impedance receiver at the host, and is terminated at the client. As described elsewhere, reverse data or data over the reverse link can be transferred or sent at data rates greater than the reciprocal of the round-trip delay in the cable. The MDDI_Stb+ and MDDI_Stb- conductors or signals are only driven by the host.

[00474] An exemplary configuration of elements useful for achieving the drivers, receivers, and terminations for transferring signals as part of the inventive MDDI are shown in FIG. 42. This exemplary interface uses low voltage sensing, here 200 mV, with less than 1 volt power swings and low power drain. The driver of each signal pair has a differential current output. While receiving MDDI packets, the MDDI_Data and MDDI_Stb pairs use a conventional differential receiver with a differential voltage threshold of zero volts. In the hibernation state the driver outputs are disabled and the parallel-termination resistors pull the differential voltage on each signal pair to zero volts. During hibernation, a special receiver on the MDDI_Data0 pair has an offset

input differential voltage threshold of positive 125 mV, which causes the hibernation line receiver to interpret the un-driven signal pair as a logic-zero level.

[00475] The differential voltage of a differential pair is defined as the difference of the voltage on the positive (+) signal minus the voltage on the negative (-) signal. The names of differential pair signals end with either “+” or “-”, which indicates the positive or negative signal of the pair, respectively. The output current of the driver of a differential pair is defined as the current flowing out of the positive (+) output. The current passing through the negative (-) output of a differential driver is always equal in magnitude but opposite in direction compared to the current passing through the positive (+) output of the same differential driver.

[00476] Sometimes the host or client simultaneously drives the differential pair to a logic-one level or a logic-zero level to guarantee a valid logic-level on the pair when the direction of data flow changes (from host-to-client or client-to-host). The output voltage range and output specifications are still met with simultaneously driven outputs driven to the same logic level. In some systems it may be necessary to drive a small current into the terminated differential pair to create a small offset voltage at certain times during hibernation and when the link is waking up from the hibernation state. In those situations, the enabled offset-current bias circuits drive the current levels referred to as: $I_{\text{ESD-and-Rx}}$ – internal ESD diode and differential receiver input with $I_{\text{ESD-and-Rx}} \leq 1 \mu\text{A}$ typically; $I_{\text{Tx-Hi-Z}}$ – differential driver output in the high-impedance state, with $I_{\text{Tx-Hi-Z}} \leq 1 \mu\text{A}$ typically; and $I_{\text{external-ESD}}$ – the leakage through the external ESD protection diodes, with $I_{\text{external-ESD}} \leq 3 \mu\text{A}$ typically.

[00477] Each of these leakage currents is illustrated in FIG. 47. The pull-up and pull-down circuits must achieve the minimum differential voltage under the worst-case leakage conditions described above when all occur simultaneously. The total leakage is $\leq 4 \mu\text{A}$ for internal mode without external ESD protection diodes and $\leq 10 \mu\text{A}$ for external mode with external ESD protection.

[00478] The electrical parameters and characteristics of the differential line drivers and line receivers are described for one exemplary embodiment in Tables IXa-IXd. Functionally, the driver transfers the logic level on the input directly to a positive output, and the inverse of the input to a negative output. The delay from input to outputs is well-matched to the differential line which is driven differentially. In most

implementations, the voltage swing on the outputs is less than the swing on the input to minimize power consumption and electromagnetic emissions. In one embodiment, there is a minimum voltage swing of around 0.5V. However, other values can be used, as would be known by those skilled in the art, and the inventors contemplate a smaller value in some embodiments, depending on design constraints.

[00479] The differential line receivers have the same characteristic as a high-speed voltage comparator. In FIG. 41, the input without the bubble is the positive input and the input with the bubble is the negative input. The output is a logic one if: $(V_{\text{input}+}) - (V_{\text{input}-})$ is greater than zero. Another way to describe this is a differential amplifier with very large (virtually infinite) gain with the output clipped at logic 0 and 1 voltage levels.

[00480] The delay skew between different pairs should be minimized to operate the differential transmission system at the highest potential speed.

Table IXa

Host Transmitter Electrical Specifications

Parameter	Description	Min	Max	Unit
$V_{\text{output-Range}}$	Allowable host driver output voltage range with respect to host ground	0.35	1.60	V
$I_{\text{OD}+}$	Driver differential output high current corresponding to a logic-one level (while driving the terminated transmission line)	2.5	4.5	mA
$I_{\text{OD}-}$	Driver differential output low current corresponding to a logic-zero level (while driving the terminated transmission line)	-4.5	-2.5	mA
$T_{\text{Rise-Fall}}$	Rise and fall time (between 20% and 80% amplitude) of driver output, measured in differential mode	425	Note 1	psec
$T_{\text{skew-pair}}$	Skew between positive and negative outputs of the same differential pair (intra-pair skew)		125	Psec
$T_{\text{Differential-Skew}}$	Peak delay skew between one differential pair and any other differential pair.		See above	Psec
T_A	Jitter, bit boundary to center crossing	0	$T_B - 283$	Psec
$T_{\text{B-TP0-DRVR}}$	Jitter, bit boundary to minimum output level	0	See above	Psec

Note 1: The maximum rise and fall time is either 30% of the interval to transmit one bit on one differential pair or 100 nsec, whichever is smaller.

Table IXb

Client Transmitter Electrical Specifications

Parameter	Description	Min	Max	Unit
V _{output-Range-Ext}	Allowable client driver output voltage range with respect to client ground (External Mode)	0	1.25	V
V _{output-Range-Int}	Allowable client driver output voltage range with respect to client ground (Internal Mode)	0.35	1.60	V
I _{OD+}	Driver differential output high current corresponding to a logic-one level (while driving the equivalent of the pull-up and pull-down circuits that exist at the host and client)	2.5	4.5	mA
I _{OD-}	Driver differential output low current corresponding to a logic-zero level (while driving the equivalent of the pull-up and pull-down circuits that exist at the host and client)	-4.5	-2.5	mA
T _{Rise-Fall}	Rise and fall time (between 20% and 80% amplitude) of driver output, measured in differential mode	425	Note 1	psec
T _{skew-pair}	Skew between positive and negative outputs of the same differential pair (intra-pair skew)		125	psec
T _{Differential-Skew}	Peak delay skew between one differential pair and any other differential pair.		See above	psec
T _A	Jitter, bit boundary to center crossing		T _B - 283	psec
T _{B-TP4-DRVr}	Jitter, bit boundary to minimum output level		See above	psec

Note 1: The maximum rise and fall time is 30% of the interval to transmit one bit on one differential pair or 100 nsec, whichever is smaller.

Table IXc

Client Receiver Electrical Specifications

Parameter	Description	Min	Typ	Max	Unit
V _{IT+}	Receiver differential input high threshold voltage. Above this differential voltage the input signal is interpreted as a logic-one level.		0	50	mV
V _{IT-}	Receiver differential input low threshold voltage. Below this differential voltage the input signal is interpreted as a logic-zero level.	-50	0		mV
V _{IT+}	Receiver differential input high threshold voltage (offset for hibernation wake-up). Above this differential voltage the input signal is interpreted as a logic-one level.		125	175	mV
V _{IT-}	Receiver differential input low	75	125		mV

	threshold voltage (offset for hibernation wake-up). Below this differential voltage the input signal is interpreted as a logic-zero level.				
$V_{\text{Input-Range}}$	Allowable receiver input voltage range with respect to client ground.	0		1.65	V
R_{term}	Parallel termination resistance value	98	100	102	Ω
I_{in}	Input leakage current	-10		10	μA
C_{pad}	Capacitance of pad to client ground (note 1)			5	pF
C_{diff}	Capacitance between the two signals of a differential pair (note 1)			1	pF
$T_{\text{skew-pair-INT}}$	Skew caused by the differential receiver between positive and negative inputs of the differential receiver of the same differential pair (intra-pair skew). Internal Mode			250	psec
$T_{\text{skew-pair-EXT}}$	Intra-pair skew, External Mode			50	psec
$T_{\text{Differential-Skew}}$	Peak delay skew between one differential pair and any other differential pair.			See above	psec
T_A	Jitter, bit boundary to center crossing			$T_B - 38.5$	psec
$T_{B\text{-}TP4\text{-}RCVR\text{-}INT}$	Jitter, bit boundary to minimum input level (Internal Mode)	0		See above	psec
$T_{B\text{-}TP4\text{-}RCVR\text{-}EXT}$	Jitter, bit boundary to minimum input level (External Mode)	0		See above	psec

Table IXd

Host Receiver Electrical Specifications

Parameter	Description	Min	Typ	Max	Unit
$V_{\text{IT}+}$	Receiver differential input high threshold voltage (non-offset). Above this differential voltage the input signal is interpreted as a logic-one level.		0	50	mV
$V_{\text{IT}-}$	Receiver differential input low threshold voltage (non-offset). Below this differential voltage the input signal is interpreted as a logic-zero level.	-50	0		mV
$V_{\text{IT}+}$	Receiver differential input high threshold voltage (offset for hibernation wake-up). Above this differential voltage the input signal is interpreted as a logic-one level.		125	175	mV
$V_{\text{IT}-}$	Receiver differential input low threshold	75	125		mV

	voltage (offset for hibernation wake-up). Below this differential voltage the input signal is interpreted as a logic-zero level				
$V_{\text{Input-Range}}$	Allowable receiver input voltage range with respect to host ground.	0		1.65	V
I_{in}	Input leakage current (excluding hibernate bias)	-10		10	μA
C_{pad}	Capacitance of pad to host ground			5	pF
C_{diff}	Capacitance between the two signals of a differential pair.			1	pF
$T_{\text{skew-pair}}$	Skew caused by the differential receiver between positive and negative inputs of the differential receiver of the same differential pair (intra-pair skew).			250	psec
$T_{\text{skew-pair-EXT}}$	Intra-pair skew, External Mode			50	psec
T_A	Jitter, bit boundary to center crossing			$T_B - 38.5$	psec
$T_{B\text{-TP0-RCVR-INT}}$	Jitter, bit boundary to minimum output level (External Mode)			See above	psec
$T_{B\text{-TP0-RCVR-EXT}}$	Jitter, bit boundary to minimum output level (External Mode)			See above	psec

[00481] In FIG. 42, a host controller 4202 and a client or display controller 4204 are shown transferring packets over the communication link 4206. The host controller employs a series of three drivers 4210, 4212, and 4214 to receive the host DATA and STB signals to be transferred, as well as to receive the client Data signals to be transferred, while the client employs the three drivers 4230, 4232, and 4234. The driver responsible for passage of the host DATA (4212) employs an enable signal input to allow activation of the communication link generally only when transfer from the host to the client is desired. Since the STB signal is formed as part of the transfer of data, no additional enable signal is employed for that driver (4212). The inputs of each of the client DATA and STB receivers (4132, 4230) have termination impedances or resistors 4218 and 4220, respectively placed across them. Driver 4234 in the client controller is used to prepare the data signals being transferred from the client to the host, where driver 4214 on the input side, processes the data.

[00482] The special receivers (drivers) 4216 and 4236 are coupled or connected to the DATA lines, and generate or use the 125 mV voltage offset previously discussed, as

part of the hibernation control discussed elsewhere. The offsets cause the hibernation line receivers to interpret un-driven signal pairs as a logic-zero level.

[00483] The above drivers and impedances can be formed as discrete components or as part of a circuit module, or an Application Specific Integrated Circuit (ASIC) which acts as a more cost effective encoder or decoder solution.

[00484] It can be easily seen that power is transferred to the client device, or display, from the host device using the signals labeled HOST_Pwr and HOST_Gnd over a pair of conductors. The HOST_Gnd portion of the signal acts as the reference ground and the power supply return path or signal for the client device. The HOST_Pwr signal acts as the client device power supply which is driven by the host device. In an exemplary configuration, for low power applications, the client device is allowed to draw up to 500 mA. The HOST_Pwr signal can be provided from portable power sources, such as but not limited to, a lithium-ion type battery or battery pack residing at the host device, and may range from 3.2 to 4.3 volts with respect to HOST_Gnd.

VII. Timing Characteristics

A. Overview

[00485] The steps and signal levels employed to enter a hibernation state (no service requested, desired, or required), and to secure service for a client from the host, either by host- or client initiation, are illustrated in FIGs. 43A, 43B, and 43C, respectively. In FIGs. 43A, 43B, and 43C, the first part of signals being illustrated shows a Link Shutdown Packet being transferred from the host and the data line is then driven to a logic zero state using the high-impedance bias circuit. No data is being transmitted by the client, or host, which has its driver disabled. A series of strobe pulses for the MDDI_Stb signal line can be seen at the bottom, since MDDI_Stb is active during the Link Shutdown Packet. Once this packet ends, the logic level changes to zero as the host drives the bias circuit and logic to zero. This represents the termination of the last signal transfer or service from the host, and could have occurred at any time in the past, and is included to show the prior cessation of service, and the state of the signals prior to service commencement. If desired, such a signal can be sent just to reset the

communication link to the proper state without a 'known' prior communication having been undertaken by this host device.

[00486] As shown in FIG. 43A, and discussed for the Link Shutdown Packet above, in the low-power hibernation state, the MDDI_Data0 driver is disabled into a high-impedance state starting after the 16th to 48th MDDI_Stb cycles or pulses after the last bit of the All Zeros field in the Link Shutdown Packet. For Type-2, Type-3, or Type-4 links the MDDI_Data1 through MDDI_DataPwr7 signals are also placed in a high-impedance state at the same time that the MDDI_Data0 driver is disabled. As described in the definition of the All Zeros field, MDDI_Stb toggles for 64 cycles (or as desired for system design) following the MSB of the CRC field of the Link Shutdown Packet to allow processing by the client to be completed and facilitate an orderly shutdown in a client controller. One cycle is a low-to-high transition followed by a high-to-low transition, or a high-to-low transition followed by a low-to-high transition. After the All Zeros field is sent, the MDDI_Stb and MDDI_Data0 drivers in the host are disabled, and the host enters the low-power hibernation state. After some period of time, the host commences the link restart sequence as shown in FIGs. 43b and 43c, by enabling the MDDI_Data0 and MDDI_Stb lines or driver outputs, and begins to toggle MDDI_Stb, as part of either at a host or client initiated wake-up request.

[00487] As shown in FIG. 43B, after some time passes with the signal output from drivers for MDDI_Data0 and MDDI_Stb disabled, a host initiates service or wake-up from hibernation by enabling its MDDI_Stb driver for a period of time designated $t_{\text{stb-data-enbl}}$, during which the line is driven to a logic zero level, until it is completely enabled and then enabling its MDDI_Data0 driver. The host holds MDDI_Stb at logic-zero level after MDDI_Data0 reaches a high or a logic one level which occurs over a period of time designated $t_{\text{client-startup}}$. At the end of the $t_{\text{client-startup}}$ period the host then toggles the MDDI_Stb signal or line. The host drives the MDDI_Data0 line high, a logic-one level, while the client does not drive MDDI_Data0, for a period designated $t_{\text{restart-high}}$ and then drives the MDDI_Data0 line low, or to logic-zero level, for a period designated $t_{\text{restart-low}}$. After this, the first forward traffic begins with a Sub-Frame Header Packet, and the forward traffic packets are then transferred. The MDDI_Stb signal is active during the $t_{\text{restart-low}}$ period and the subsequent Sub-Frame Header Packet.

[00488] As shown in FIG. 43C, after some time passes with the signal output from drivers for MDDI_Data0 and MDDI_Stb disabled, a client initiates a service request or

wake-up from hibernation by enabling an offset in the MDDI_Stb receiver or output signal for a period of time designated $t_{\text{stb-data-enb}}$, as discussed above, before the host enables its MDDI_Stb driver. The client then enables its MDDI_Data0 driver for a period of time designated $t_{\text{host-detect}}$, during which the line is driven to a logic zero level, before the host begins MDDI_Stb toggling.

[00489] A certain amount of time passes or may be needed before the host detects the request during $t_{\text{host-detect}}$, after which the host responds by holding MDDI_Stb at logic-zero level for the period designated $t_{\text{stb-startup}}$ before the host begins toggling MDDI_Stb with a link startup sequence by driving the MDDI_Data0 to a logic-one or high level during the $t_{\text{restart-high}}$ period. When the client recognizes the first pulse on MDDI_Stb it disables the offset in its MDDI_Stb receiver. The client continues to drive MDDI_Data0 to a logic-one level or a period designated $t_{\text{client-detect}}$ until it detects the host driving the line. At this point, the client de-asserts the request, and disables its MDDI_Data0 driver so that the output from the client goes to a logic-zero level again, and the host is driving MDDI_Data0. As before, the host continues to drive MDDI_Data0 to a logic-one level for the $t_{\text{restart-high}}$ period, and then drives the MDDI_Data0 line low for the $t_{\text{restart-low}}$ period, after which the first forward traffic begins with a Sub-Frame Header Packet. The MDDI_Stb signal is active during the $t_{\text{restart-low}}$ period and the subsequent Sub-Frame Header Packet.

[00490] Table X shows representative times or processing periods for the length of the various periods discussed above, and the relationship to exemplary minimum and maximum data rates, where:

$$t_{\text{bit}} = \frac{1}{\text{Link_Data_Rate}}, \text{ where Link_Data_Rate is the bit rate of a single data}$$

pair.

Table X

Parameter	Description	Min.	Typ.	Max.	Units
$1/t_{BIT-min-perf}$	Link data rate for a minimum performance device	0.001		1.1	Mbps
$1/t_{BIT-max-perf}$	Maximum link data rate range for a device, external	0.001		400	Mbps
$1/t_{BIT-max-perf}$	Maximum link data rate range for a device, internal	0.001		550	Mbps
	Reverse Link data rate	0.0005		50	Mbps
t_{BIT}	Period of one forward link data bit, external mode	2.5		10^0	nsec
t_{BIT}	Period of one forward link data bit, internal mode	1.8		10^0	nsec
$t_{restart-high}$	Duration of host link restart high pulse	140	150	160	Stb clks
$t_{restart-low}$	Duration of host link restart low pulse	50	50	50	Stb clks
$t_{stb-data-enabl}$	MDDI_Stb completely enabled to MDDI_Data0 enabled link restart sequence	0			μ sec
$t_{client-startup}$	Time for host to hold MDDI_Stb at logic-zero level after MDDI_Data0 reaches logic-high level	200			nsec
$t_{host-detect}$	Time from MDDI_Data0 high to MDDI_Stb toggling	0		1000	μ sec
$t_{client-detect}$	Time for client to detect MDDI_Data0 at logic-high level performance device	60		80	Stb clks
$t_{stb-startup}$	Time for host to hold MDDI_Stb at logic-zero level before host begins toggling MDDI_Stb	200			nsec

[00491] Those skilled in the art will readily understand that the functions of the individual elements illustrated in FIGs. 41 and 42, are well known, and the function of the elements in FIG. 42 is confirmed by the timing diagram in FIGS 43a, 43b, and 43c. Details about the series terminations and hibernation resistors that are shown in FIG. 42 were omitted from FIG. 41 because that information is unnecessary for a description of how to perform the Data-Strobe encoding and recover the clock from it.

B. Data-Strobe Timing Forward Link

[00492] The switching characteristics for the transfer of data on the forward link from the host driver output is shown in Table XI-1. Table XI presents a tabular form of the minimum and maximum desired, versus typical times for certain signal transitions to occur. For example, the typical length of time for a transition to occur from the start to the end of a data value (output of a '0' or '1'), a Data0 to Data0 transition, termed t_{dd} -(host-output), is t_{bit} while the minimum time is about $t_{bit}-0.5$ nsec., and the maximum is about $t_{bit}+0.5$ nsec. The relative spacing between transitions on the Data0, other data lines (DataX), and the strobe lines (Stb), is illustrated in FIG. 44 where the Data0 to Strobe, Strobe to Strobe, Strobe to Data0, Data0 to non-Data0, non-Data0 to non-Data0, non-Data0 to Strobe, and Strobe to non-Data0 transitions are shown, which are referred to as t_{ds} -(host-output), t_{ss} -(host-output), t_{sd} -(host-output), t_{ddx} -(host-output), t_{ddx} -(host-output), t_{dxs} -(host-output), and t_{sdx} -(host-output), respectively.

Table XI-1

Parameter	Description	Min.	Typ.	Max.	Units
t_{dd} -(host-output)	Data0 to Data0 transition	$t_{bit} - 0.5$	t_{bit}	$t_{bit} + 0.5$	nsec
t_{ds} -(host-output)	Data0 to Strobe transition	$t_{bit} - 0.8$	t_{bit}	$t_{bit} + 0.8$	nsec
t_{ss} -(host-output)	Strobe to Strobe transition	$t_{bit} - 0.5$	t_{bit}	$t_{bit} + 0.5$	nsec
t_{sd} -(host-output)	Strobe to Data0 transition	$t_{bit} - 0.8$	t_{bit}	$t_{bit} + 0.8$	nsec
t_{ddx} -(host-output)	Data0 to non-Data0 transition		t_{bit}		nsec
t_{ddx} -(host-output)	non-Data0 to non-Data0 transition	$t_{bit} - 0.5$	t_{bit}	$t_{bit} + 0.5$	nsec
t_{dxs} -(host-output)	non-Data0 to Strobe transition		t_{bit}		nsec
t_{sdx} -(host-output)	Strobe to non-Data0 transition		t_{bit}		nsec

[00493] The typical MDDI timing requirements for the client receiver input for the same signals transferring data on the forward link is shown in Table XI-2. Since the same signals are being discussed but time delayed, no new figure is needed to illustrate the signal characteristics or meaning of the respective labels, as would be understood by those skilled in the art.

Table XI-2

Parameter	Description	Min.	Typ.	Max.	Units
$t_{dd}(\text{client-input})$	Data0 to Data0 transition	$t_{bit} - 1.0$	t_{bit}	$t_{bit} + 1.0$	nsec
$t_{ds}(\text{client-input})$	Data0 to Strobe transition	$t_{bit} - 1.5$	t_{bit}	$t_{bit} + 1.5$	nsec
$t_{ss}(\text{client-input})$	Strobe to Strobe transition	$t_{bit} - 1.0$	t_{bit}	$t_{bit} + 1.0$	nsec
$t_{sd}(\text{client-input})$	Strobe to Data0 transition	$t_{bit} - 1.5$	t_{bit}	$t_{bit} + 1.5$	nsec
$t_{ddx}(\text{host-output})$	Data0 to non-Data0 transition		t_{bit}		nsec
$t_{dndx}(\text{host-output})$	non-Data0 to non-Data0 transition		t_{bit}		nsec
$t_{dxx}(\text{host-output})$	non-Data0 to Strobe transition		t_{bit}		nsec
$t_{sdx}(\text{host-output})$	Strobe to non-Data0 transition		t_{bit}		nsec

[00494] FIGs. 45 and 46 illustrate the presence of a delay in response that can occur when the host disables or enables the host driver, respectively. In the case of a host forwarding certain packets, such as the Reverse Link Encapsulation Packet or the Round Trip Delay Measurement Packet, the host disables the line driver after the desired packets are forwarded, such as the Parameter CRC, Strobe Alignment, and All Zero packets illustrated in FIG. 45 as having been transferred. However, as shown in FIG. 45, the state of the line does not necessarily switch from '0' to a desired higher value instantaneously, although this is potentially achievable with certain control or circuit elements present, but takes a period of time termed the host Driver Disable Delay period to respond. While it could occur virtually instantly such that this time period is 0 nanoseconds (nsec.) in length, it could more readily extend over some longer period with 10 nsec. being a desired maximum period length, which occurs during the Guard Time 1 or Turn Around 1 packet periods.

[00495] Looking in FIG. 46, one sees the signal level change undergone when the host Driver is enabled for transferring a packet such as the Reverse Link Encapsulation Packet or the Round Trip Delay Measurement Packet. Here, after the Guard Time 2 or Turn Around 2 packet periods, the host driver is enabled and begins to drive a level, here '0', which value is approached or reached over a period of time termed the Host Driver Enable Delay period, which occurs during the Driver Re-enable period, prior to the first packet being sent.

[00496] A similar process occurs for the drivers and signal transfers for the client device, here a display. The general guidelines for the length of these periods, and their respective relationships are shown in XII, below.

Table XII

Description	Min.	Max.	Units
Host Driver Disable Delay	0	10	nsec
Host Driver Enable Delay	0	2.0	nsec
Display Driver Disable Delay	0	10	nsec
Display Driver Enable Delay	0	2.0	nsec

C. Host And Client Output Enable And Disable Times

[00497] The switching characteristics and relative timing relationships for Host and Client output enabled and disable time or operations relative to the Reverse Link Encapsulation Packet structure and period, is shown in FIG. 48. The driver output functions or operations are labeled as: $t_{\text{host-enable}}$ for the Host output enable time, $t_{\text{host-disable}}$ for the Host output disable time, $t_{\text{client-enable}}$ for the Client output enable time, and $t_{\text{client-disable}}$ for the Client output disable time. Typical times for certain signal transitions are discussed below. The minimum period for these operations would be zero nanoseconds, with typical or maximum values determined from the system design employing the interface, possibly on the order of 8 nanoseconds, or more.

[00498] The general guidelines for the length of these periods, (host and client enable/disable times) and their respective relationships are shown in XIII, below.

Table XIII

Parameter	Description	Min.	Typ.	Max.	Units
$t_{\text{host-enable}}$	Host output enable time	0		$24 \cdot t_{\text{BIT}}$	nsec
$t_{\text{host-disable}}$	Host output disable time, entire length of the Turn-Around 1 field	0		$24 \cdot t_{\text{BIT}}$	nsec
$t_{\text{client-enable}}$	Client output enable time, entire length of the Turn-Around 1 field	0		$24 \cdot t_{\text{BIT}}$	nsec
$t_{\text{client-disable}}$	Client output disable time, measured from the end of the last bit of the Turn-Around 2 field	0		$24 \cdot t_{\text{BIT}}$	nsec

VIII. Implementation of Link Control (Link Controller Operation)

A. State Machine Packet Processor

[00499] Packets being transferred over a MDDI link are dispatched very rapidly, typically at a rate on the order of 300 Mbps or more, such as 400 Mbps, although lower rates are certainly accommodated, as desired. This type of bus or transfer link speed is too great for currently commercially available (economical) general-purpose microprocessors or the like to control. Therefore, a practical implementation to accomplish this type of signal transfer is to utilize a programmable state machine to parse the input packet stream to produce packets that are transferred or redirected to the appropriate audio-visual subsystem for which they are intended. Such devices are well known and use circuits generally dedicated to a limited number of operations, functions, or states to achieve a desired high speed or very high speed operation.

[00500] General purpose controllers, processors, or processing elements, can be used to more appropriately act upon or manipulate some information such as control or status packets, which have lower speed demands. When those packets (control, status, or other pre-defined packets) are received, the state machine should pass them through a data buffer or similar processing element to the general-purpose processor so the packets can be acted upon to provide a desired result (effect) while the audio and visual packets are transferred to their appropriate destination for action. If future, microprocessors or other general purpose controllers, processors, or processing elements are manufactured to achieve higher data rate processing capabilities, then the states or state machine discussed below might also be implemented using software control of such devices, typically as programs stored on a storage element or media.

[00501] The general purpose processor function can be realized in some embodiments by taking advantage of the processing power, or excess cycles available for, microprocessors (CPUs) in computer applications, or controllers, processors, digital signal processors (DSPs), specialized circuits, or ASICs found in wireless devices, in much the same manner as some modems or graphics processors utilize the processing power of CPUs found in computers to perform some functions and reduce hardware complexity and costs. However, this cycle sharing or usage can negatively impact the processing speed, timing, or overall operation of such elements, so in many applications, dedicated circuits or elements are preferred for this general processing.

[00502] In order for image data to be viewed on a display (micro-display), or to reliably receive all packets sent by the host device, the client signal processing is synchronized with the forward link channel timing. That is, signals arriving at the client and the client circuits need to be substantially time synchronized for proper signal processing to occur. A high level diagram of states achieved by signal for one embodiment is presented in the illustration of FIG. 49. In FIG. 49, the possible forward link synchronization "states" for a state machine 4900 are shown being categorized as one ASYNC FRAMES STATE 4904, two ACQUIRING SYNC STATES 4902 and 4906, and three IN-SYNC STATES 4908, 4910, and 4912.

[00503] As shown by starting step or state 4902, the display or client, such as a presentation device, starts in a pre-selected "no sync" state, and searches for a unique word in the first sub-frame header packet that is detected. It is to be noted that this no sync state represents the minimum communication setting or "fall-back" setting in which a Type 1 interface is selected. When the unique word is found during the search, the client saves the sub-frame length field. There is no checking of the CRC bits for processing on this first frame, or until synchronization is obtained. If this sub-frame length is zero, then sync state processing proceeds accordingly to a state 4904 labeled here as the "async frames" state, which indicates that synchronization has not yet been achieved. This step in the processing is labeled as having encountered cond 3, or condition 3, in FIG. 49. Otherwise, if the frame length is greater than zero, then the sync state processing proceeds to a state 4906 where the interface state is set as "found one sync frame." This step in the processing is labeled as encountering cond 5, or condition 5, in FIG. 49. In addition, if the state machine sees a frame header packet and good CRC determination for a frame length greater than zero, processing proceeds to the "found one sync frame" state. This is labeled as meeting cond 6, or condition 6, in FIG. 49.

[00504] In each situation in which the system is in a state other than "no sync," if a packet with a good CRC result is detected, then the interface state is changed to the "in-sync" state 4908. This step in the processing is labeled as having encountered cond 1, or condition 1, in FIG. 49. On the other hand, if the CRC in any packet is not correct, then the sync state processing proceeds or returns to the interface state 4902 of "NO SYNC FRAME" state. This portion of the processing is labeled as encountering cond 2, or condition 2, in the state diagram of FIG. 49.

B. Acquisition Time for Sync

- [00505] The interface can be configured to accommodate a certain number of "sync errors" prior to deciding that synchronization is lost and returning to the "NO SYNC FRAME" state. In FIG. 49, once the state machine has reached the "IN-SYNC STATE" and no errors are found, it is continuously encountering a cond 1 result, and remains in the "IN-SYNC" state. However once one cond 2 result is detected, processing changes the state to a "one-sync-error" state 4910. At this point, if processing results in detecting another cond 1 result, then the state machine returns to the "in-sync" state, otherwise it encounters another cond 2 result, and moves to a "TWO-SYNC-ERRORS" state 4912. Again, if a cond 1 occurs, processing returns the state machine to the "IN-SYNC" state. Otherwise, another cond 2 is encountered and the state machine returns to the "no-sync" state. It is also understandable that should the interface encounter a "link shutdown packet," then this will cause the link to terminate data transfers and return to the "no-sync frame" state as there is nothing to synchronize with, which is referred to as meeting cond 4, or condition 4, in the state diagram of FIG. 49.
- [00506] It is understood that it is possible for there to be a repeating "false copy" of the unique word which may appear at some fixed location within the sub-frame. In that situation, it is highly unlikely that the state machine will synchronize to the sub-frame because the CRC on the sub-frame Header Packet must also be valid when processed in order for the MDDI processing to proceed to the "IN SYNC" state.
- [00507] The sub-frame length in the sub-frame Header Packet may be set to zero to indicate that the host will transmit only one sub-frame before the link is shut down, and the MDDI is placed in or configured into an idle hibernation state. In this case, the client must immediately receive packets over the forward link after detecting the sub-frame Header Packet because only a single sub-frame is sent before the link transitions to the idle state. In normal or typical operations, the sub-frame length is non-zero and the client only processes forward link packets while the interface is in those states collectively shown as "IN-SYNC" states in FIG. 49.
- [00508] An external mode client device may be attached to the host while the host is already transmitting a forward link data sequence. In this situation, the client must synchronize to the host. The time required for a client to synchronize to the forward link signal is variable depending on the sub-frame size and the forward link data rate.

The likelihood of detecting a “false copy” of the unique word as part of the random, or more random, data in the forward link is greater when the sub-frame size is larger. At the same time, the ability to recover from a false detection is lower, and the time taken to do so is longer, when a forward link data rate is slower.

[00509] For one or more embodiments, it is recommended or understood that a MDDI host should perform certain additional steps to ensure that the MDDI reverse link is stable before it stops forward link transmission to go to a low power mode or to shut down the link completely.

[00510] One problem that can occur is that if a host uses an incorrect measurement of the round-trip delay value this can cause all subsequently received reverse data transmission from the client to fail even though the forward link appears to be fine. This could happen if the host tries to send a Round Trip Delay Measurement Packet when the client is not in sync with the forward link, or due to an extreme ambient temperature change that causes a corresponding large change in the propagation delay of the differential drivers and receivers which affects the round trip delay. An intermittent cable or connector contact failure could also cause the client to temporarily lose synchronization and then regain sync, during which time, it may miss receiving a Round Trip Delay Measurement Packet. Subsequent reverse link packets would not be able to be decoded properly by the host.

[00511] Another type of problem that can occur is that if the client temporarily loses sync and the host sends a Link Shutdown Packet before the client is able to regain sync. The host will be in hibernation while the client is unable to enter the hibernation state because it did not receive the Link Shutdown Packet and does not have a clock because the link is in hibernation.

[00512] One technique or embodiment useful for overcoming such problems is to have the host ensure that the client is in sync with the forward link before putting the link into the hibernation state. If the MDDI host is unable to do this or does not have such an opportunity, such as when it loses power or the link is abruptly broken or fails due to a cable, conductor, or connector separation, break, or disconnection occurring during operation, then the host should first try to ensure that the client is in sync before starting a round-trip delay measurement process or sending a Reverse Link Encapsulation packet.

[00513] A host can observe the CRC Error Count field in a Client Request and Status packet sent by the client to determine the forward link integrity. This packet is requested by the host from the client. However, in the event of a major link failure or disruption, this request will most likely go unanswered since a client will not be able to properly decode the packet, or maybe even receive it altogether. The request for the CRC Error Count using the Client Request and Status Packet sent in a Reverse Link Encapsulation Packet acts as a first integrity check, a sort of first line of defense. In addition, a host can send a Round Trip Delay Measurement Packet to confirm whether or not the assumption about the client having fallen out of sync is a valid one or not. If the client does not respond to a Round Trip Delay Measurement Packet, the host will conclude that the client is out of sync and can then start the process of getting it back in sync.

[00514] Once the host concludes that the client has more than likely lost synchronization with the forward link, it waits until the next subframe header before attempting to send any packets other than filler packets. This is done in order to allow a client enough time to detect or look for one unique word contained in the subframe header packet. Following this, the host may assume that the client would have reset itself since it would not have found the unique word at the correct location. At this point, the host may follow the subframe header packet with a Round Trip Delay Measurement Packet. If the client still does not respond correctly to the Round Trip Delay Measurement Packet, the host may repeat the resynchronization process. A correct response is one in which the client sends the specified sequence back to the host in the Measurement Period of the Round Trip Delay Measurement Packet. If this sequence is not received, then attempts to receive reverse data in a Reverse Link Encapsulation Packet will fail. Continued failure of this nature may indicate some other system error which will have to be addressed in some other manner, and is not part of the link synchronization at this point.

[00515] However, if after a successful Round Trip Delay Measurement Packet the host still sees corrupted data or no response in the Reverse Link Encapsulation Packets, it should confirm the reverse data sampling is correct by re-sending a Round Trip Delay Measurement Packet. If this is not successful after a number of attempts it is recommended for one embodiment that the host reduce the reverse data rate by increasing the reverse rate divisor value.

[00516] The host should perform the Link Failure Detection and possibly the Link Resynchronization steps described above before placing the MDDI link into the hibernation state. This will generally ensure that the Round Trip Delay Measurement Packet performed when the link is restarted later on is successful. If the host has no reason to suspect a link failure, and a correct response to a Reverse Link Encapsulation Packet and zero forward link CRC errors is being reported by the client, a host may assume that everything is operating or functioning accordingly or appropriately (no link failure for example) and proceed with the power down/hibernation process.

[00517] Another manner in which a host can test for synchronization is for the host to send the Round Trip Delay Measurement Packet and confirm the proper response from the client. If the proper response is received by the host, it can reasonably be assumed that the client is successfully interpreting forward link packets.

C. Initialization

[00518] As stated earlier, at the time of "start-up," the host configures the forward link to operate at or below a minimum required, or desired, data rate of 1 Mbps, and configures the sub-frame length and media-frame rate appropriately for a given application. That is, both the forward and reverse links begin operation using the Type 1 interface. These parameters are generally only going to be used temporarily while the host determines the capability or desired configuration for the client display (or other type of client device). The host sends or transfers a sub-frame Header Packet over the forward link followed by a Reverse Link Encapsulation Packet which has bit '0' of the Request Flags set to a value of one (1), in order to request that the display or client responds with a Client Capability Packet. Once the display acquires synchronization on (or with) the forward link, it sends a Client Capability Packet and a Client Request and Status Packet over the reverse link or channel.

[00519] The host examines the contents of the Client Capability Packet in order to determine how to reconfigure the link for optimal or a desired level of performance. The host examines the Protocol Version and Minimum Protocol Version fields to confirm that the host and client use versions of the protocol that are compatible with each other. The protocol versions generally remain as the first two parameters of the client capability Packet so that compatibility can be determined even when other

elements of the protocol might not be compatible or completely understood as being compatible.

[00520] In internal mode the host can know the parameters of the client in advance without having to receive a Client Capability Packet. The link may start up at any data rate at which the host and client can both operate. In many embodiments, a system designer will most likely choose to start the link at the maximum achievable data rate to hasten data transfer, however, this is not required and may not be used in many situations. For internal mode operation, the frequency of the strobe pulses used during the link restart from hibernation sequence will usually be consistent with this desired rate.

D. CRC Processing

[00521] For all packet types, the packet processor state machine ensures that the CRC checker is controlled appropriately or properly. It also increments a CRC error counter when a CRC comparison results in one or more errors being detected, and it resets the CRC counter at the beginning of each sub-frame being processed.

E. Alternative Loss Of Synchronization Check

[00522] While the above series of steps or states work to produce higher data rates or throughput speed, Applicants have discovered that an alternative arrangement or change in the conditions the client uses to declare that there is a loss of synchronization with the host, can be used effectively to achieve even higher data rates or throughput. The new inventive embodiment has the same basic structure, but with the conditions for changing states changed. Additionally a new counter is implemented to aid in making checks for sub-frame synchronization. These steps and conditions are presented relative to FIG. 63, which illustrates a series of states and conditions useful in establishing the operations of the method or state machine. Only the "ACQUIRING-SYNC STATES" and "IN-SYNC STATES" portions are shown for clarity. In addition, since the resulting states are substantially the same, as is the state machine itself, they use the same numbering. However, the conditions for changing states (and the state machine operation) vary somewhat, so that all are renumbered for clarity between the two figures (1, 2, 3, 4, 5, and 6, versus 61, 62, 63, 64, and 65), as a convenience in identifying

differences. Since the ASYNC FRAME state is not considered in this discussion, one state (4904) and condition (6) are no longer used in the figure.

[00523] In FIG. 63, the system or client (for display or presentation) starts with state machine 5000 in the pre-selected "no sync" state 4902, as in FIG. 49. The first state change for changing states from the no-sync condition 4902 is in condition 64 which is the discovery of the sync pattern. Assuming that the CRC of the sub-frame header also passes on this packet (meets condition 61), the state of the packet processor state machine can be changed to the in-sync state 4908. A sync error, condition 62, will cause the state machine to shift to state 4910, and a second occurrence to state 4912. However, it has been discovered that any CRC failure of an MDDI packet will cause the state machine to move out of in-sync state 4908, to the one sync error state 4910. Another CRC failure of any MDDI packet will cause a move to the two sync failure state 4912. A packet decoded with a correct CRC value will cause the state machine to return to the in-sync state 4908.

[00524] What has been changed is to utilize the CRC value or determination for 'every' packet. That is, to have the state machine look at the CRC value for every packet to determine a loss of synchronization instead of just observing sub-frame header packets. In this configuration or process, a loss of synchronization is not determined using the unique word and just sub-frame header CRC values.

[00525] This new interface implementation allows the MDDI link to recognize synchronization failures much more quickly, and therefore, to recover from them more quickly, as well.

[00526] To make this system more robust, the client should also add or utilize a sub-frame counter. The client then checks for the presence of the unique word at the time it is expected to arrive or occur in a signal. If the unique word does not occur at the correct time, the client can recognize that a synchronization failure has occurred much more quickly than if it had to wait several (here three) packet times or periods that were greater than a sub-frame length. If the test for the unique word indicates it is not present, in other words that the timing is incorrect, then the client can immediately declare a link loss of synchronization and move to the no-sync state. The process of checking for the proper unique word presence, adds a condition 65 (cond 65) to the state machine saying that the unique word is incorrect. If a sub-frame packet is expected to be received on the client and doesn't match up, the client can immediately go to the no-

sync state 4902, saving additional time waiting for multiple sync errors (condition 62) normally encountered traversing through states 4910 and 4912.

[00527] This change uses an additional counter or counting function in the client core to count sub-frame length. In one embodiment, a count down function is used and the transfer of any packet that was currently being processed is interrupted to check for the sub-frame unique word if the counter has expired. Alternatively, the counter can count up, with the count being compared to a desired maximum or particular desired value, at which point the current packet is checked. This process protects the client from decoding packets that are incorrectly received on the client with extraordinarily long packet lengths. If the sub-frame length counter needed to interrupt some other packet that was being decoded, a loss of synchronization can be determined since no packet should cross a sub-frame boundary.

IX. Packet Processing

[00528] For each type of packet discussed above that the state machine receives, it undertakes a particular processing step or series of steps to implement operation of the interface. Forward link packets are generally processed according to the exemplary processing listed in Table XIV below.

Table XIV

<i>Packet type</i>	<i>Packet processor state machine response</i>
Sub-Frame Header (SH)	Confirms good packet, captures sub-frame length field, and sends packet parameters to a general purpose processor.
Filler (F)	Ignores data.
Video Stream (VS)	Interprets the Video Data Format Descriptor and other parameters, unpacks packed pixel data when necessary, translates pixels through the color map if necessary, and writes pixel data to appropriate locations in the bitmap.
Audio Stream (AS)	Sends audio sample rate setting to

<i>Packet type</i>	<i>Packet processor state machine response</i>
	audio sample clock generator, separates audio samples of specified size, unpacks audio sample data when necessary, and routes audio samples to appropriate audio sample FIFO
Color Map (CM)	Reads color map size and offset parameters, and writes the color map data to a color map memory or storage location.
Reverse Link Encapsulation (REL)	Facilitates sending packets in reverse direction at the appropriate time. Reverse link flags are examined, and Client Capability packets are sent as necessary. Client Request and Status packets are also sent as appropriate.
Client Capability (CC)	Sends this type of packet when requested by a host using the reverse link flags field of the Reverse Link Encapsulation Packet.
Keyboard (K)	Passes these packets to and from a general purpose processor that communicates with a keyboard type device, if one is present, and use is desired.
Pointing Device (PD)	Passes these packets to and from a general purpose processor that communicates with a pointing type device, if one is present, and use is desired.
Link Shutdown (LS)	Records fact link is shut down and informs a general-purpose processor.
Client Service Request and Status (CSRS)	Sends this packet as the first packet in the Reverse Link Encapsulation packet.
Bit Block Transfer (BPT)	Interprets packet parameters, such as Video Data Format Descriptor, determines which pixels to move first, and moves pixels in bitmap as required.
Bitmap Area Fill (BAF)	Interprets packet parameters, translates pixels through color map if necessary, and writes pixel data to appropriate locations in bitmap,

<i>Packet type</i>	<i>Packet processor state machine response</i>
Bitmap Pattern Fill (BPF)	Interprets packet parameters, unpacks packed pixel data if necessary, translates pixels through color map if necessary, and writes pixel data to appropriate locations in bitmap.
Communication Link Channel (CLC)	Sends this data directly to a general-purpose processor.
Client Service Request (CSR) during hibernation	General-purpose processor controls the low-level functions of sending request and detects contention with Link restarting on its own.
Interface Type Handoff Request (ITHR) and Interface Type Acknowledge (ITA)	May pass these packets to and from the general-purpose processor. The logic to receive this type of packet and formulate a response with an acknowledgment is substantially minimal. Therefore, this operation could also be implemented within the packet processor state machine. The resulting handoff occurs as a low-level physical layer action and is not likely to affect the functionality or functioning of the general-purpose processor.
Perform Type Handoff (PTH)	May act on such packets either directly or by transferring them to the general-purpose processor, also commanding hardware to undergo a mode change.

X. Reducing the Reverse Link Data Rate

[00529] It has been observed by the inventors that certain parameters used for the host link controller can be adjusted or configured in a certain manner in order to achieve a maximum or more optimized (scale) reverse link data rate, which is very desirable. For example, during the time used to transfer the Reverse Data Packets field of the Reverse Link Encapsulation Packet, the MDDI_Stb signal pair toggles to create a periodic data clock at half the forward link data rate. This occurs because the host link controller generates the MDDI_Stb signal that corresponds to the MDDI_Data0 signal as if it were sending all zeroes. The MDDI_Stb signal is transferred from the host to a client where it is used to generate a clock signal for transferring reverse link data from the client, with which reverse data is sent back to the host. An illustration of typical amounts of

delay encountered for the signal transfer and processing on the forward and reverse paths in a system employing the MDDI, is shown in FIG. 50. In FIG. 50, a series of delay values 1.5 nsec., 8.0 nsec., 2.5 nsec., 2.0 nsec., 1.0 nsec., 1.5 nsec., 8.0 nsec., and 2.5 nsec., are shown near processing portions for the Stb+/- generation, cable transfer-to-client, client receiver, clock generation, signal clocking, Data0+/- generation, cable transfer-to-host, and host receiver stages, respectively.

[00530] Depending on the forward link data rate and signal processing delays encountered, it may require more time than one cycle on the MDDI_Stb signal for this "round trip" effect or set of events to be completed, which results in the consumption undesirable amounts of time or cycles. To circumvent this problem, the Reverse Rate Divisor makes it possible for one bit time on the reverse link to span multiple cycles of the MDDI_Stb signal. This means that the reverse link data rate is less than the forward link rate.

[00531] It should be noted that the actual length of signal delays through the interface may differ depending on each specific host-client system or hardware being used. Although not required, each system can generally be made to perform better by using the Round Trip Delay Measurement Packet to measure the actual delay in a system so that the Reverse Rate Divisor can be set to an optimum value. The host may support either basic data sampling which is simpler but operates at a slower speed or advanced data sampling that is more complex but supports higher reverse data rates. The client capability to support both methods is considered the same

[00532] A round-trip delay is measured by having the host send a Round Trip Delay Measurement Packet to the client. The client responds to this packet by sending a sequence of ones back to the host inside of, or during, a pre-selected measurement window in that packet called the Measurement Period field. The detailed timing of this measurement was described previously. The round-trip delay is used to determine the rate at which the reverse link data can be safely sampled.

[00533] The round-trip delay measurement consists of determining, detecting, or counting the number of forward link data clock intervals occurring between the beginning of the Measurement Period field and the beginning of the time period when the 0xff, 0xff, 0x00 response sequence is received back at the host from the client. Note that it is possible that the response from the client could be received a small fraction of a forward link clock period before the measurement count was about to increment. If this

unmodified value is used to calculate the Reverse Rate Divisor it could cause bit errors on the reverse link due to unreliable data sampling. An example of this situation is illustrated in FIG. 51, where signals representing MDDI_Data at host, MDDI_Stb at host, forward link data clock inside the host, and a Delay Count are then do write the resulting pixel to the destination pixel location illustrated in graphical form. In FIG. 51, the response sequence was received from the client a fraction of a forward link clock period before the Delay Count was about to increment from 6 to 7. If the delay is assumed to be 6, then the host will sample the reverse data just after a bit transition or possibly in the middle of a bit transition. This could result in erroneous sampling at the host. For this reason, the measured delay should typically be incremented by one before it is used to calculate the Reverse Rate Divisor.

- [00534] The Reverse Rate Divisor is the number of MDDI_Stb cycles the host should wait before sampling the reverse link data. Since MDDI_Stb is cycled at a rate that is one half of the forward link rate, the corrected round-trip delay measurement needs to be divided by 2 and then rounded up to the next integer. Expressed as a formula, this relationship is:

$$reverse_rate_divisor = RoundUpToNextInteger\left(\frac{round_trip_delay + 1}{2}\right)$$

For the example given, this becomes:

$$reverse_rate_divisor = RoundUpToNextInteger\left(\frac{6 + 1}{2}\right) = 4$$

- [00535] If the round trip delay measurement used in this example were 7 as opposed to 6, then the Reverse Rate Divisor would also be equal to 4.

- [00536] The reverse link data is sampled by the host on the rising edge of the Reverse Link Clock. There is a counter or similar known circuit or device present in both the host and client (display) to generate the Reverse Link Clock. The counters are initialized so that the first rising edge of the Reverse Link Clock occurs at the beginning of the first bit in the Reverse Link Packets field of the Reverse Link Encapsulation packet. This is illustrated, for the example given below, in FIG. 52A. The counters

increment at each rising edge of the MDDI_Stb signal, and the number of counts occurring until they wrap around is set by the Reverse Rate Divisor parameter in the Reverse Link Encapsulation Packet. Since the MDDI_Stb signal toggles at one half of the forward link rate, then the reverse link rate is one half of the forward link rate divided by the Reverse Rate Divisor. For example, if the forward link rate is 200 Mbps and the Reverse Rate Divisor is 4 then the reverse link data rate is expressed as:

$$\frac{1}{2} \cdot \frac{200Mbps}{4} = 25Mbps$$

[00537] An example showing the timing of the MDDI_Data0 and MDDI_Stb signal lines in a Reverse Link Encapsulation Packet is shown in FIG. 52, where the packet parameters used for illustration have the values:

Packet Length = 1024 (0x0400)	Turn Around 1 Length = 1
Packet Type = 65 (0x41)	Turn Around 2 Length = 1
Reverse Link Flags = 0	Reverse Rate Divisor = 2
Parameter CRC = 0xdb43	All Zero is 0x00

Packet data between the Packet Length and Parameter CRC fields is:

0x00, 0x04, 0x41, 0x00, 0x02, 0x01, 0x01, 0x43, 0xdb, 0x00, ...

[00538] The first reverse link packet returned from the client is the Client Request and Status Packet having a Packet Length of 7 and a packet type of 70. This packet begins with the byte values 0x07, 0x00, 0x46, ... and so forth. However, only the first byte (0x07) is visible in FIG. 52. This first reverse link packet is time-shifted by nearly one reverse link clock period in the figure to illustrate an actual reverse link delay. An ideal waveform with zero host to client round-trip delay is shown as a dotted-line trace.

[00539] The MS byte of the Parameter CRC field is transferred, preceded by packet type, then the all zero field. The strobe from the host is switching from one to zero and back to one as the data from the host changes level, forming wider pulses. As the data goes to zero, the strobe switches at the higher rate, only the change in data on the data line causes a change near the end of the alignment field. The strobe switches at the higher rate for the remainder of the figure due to the fixed 0 or 1 levels of the data signal for extended periods of time, and the transitions falling on the pulse pattern (edge).

[00540] The reverse link clock for the host is at zero until the end of the Turn Around 1 period, when the clock is started to accommodate the reverse link packets. The arrows in the lower portion of the figure indicate when the data is sampled, as would be apparent from the remainder of the disclosure. The first byte of the packet field being transferred (here 11000000) is shown commencing after Turn Around 1, and the line level has stabilized from the host driver being disabled. The delay in the passage of the first bit, and as seen for bit three, can be seen in the dotted lines for the Data signal.

[00541] In FIG. 53, one can observe typical values of the Reverse Rate Divisor based on the forward link data rate. The actual Reverse Rate Divisor is determined as a result of a round-trip link measurement to guarantee proper reverse link operation. A first region 5302 corresponds to an area of safe operation, a second region 5304 corresponds to an area of marginal performance, while a third region 5306 indicates settings that are unlikely to function properly.

[00542] The round-trip delay measurement and Reverse Rate Divisor setting are the same while operating with any of the Interface Type settings on either the forward or reverse link because they are expressed and operated on in terms of units of actual clock periods rather than numbers of bits transmitted or received.

[00543] Typically, the largest possible Reverse Rate Divisor is half the number of bits that can be sent in the measurement window of the Round Trip Delay Measurement Packet using a Type-I interface, or for this example:

$$\frac{(512 \text{ bytes} \cdot 8 \text{ bits / byte})}{2} = 2048$$

[00544] An advanced reverse data sampling method can also be employed as an alternative that allows the reverse bit time to be smaller than the round-trip delay. For this technique a host not only measures the round-trip delay, but can also determine the phase of the response from the client with respect to an 'ideal' bit boundary of a client and link with zero delay. By knowing the phase of the client device response, a host can determine a relatively safe time to sample the reverse data bits from the client. The round-trip delay measurement indicates to a host the location of the first bit of reverse data with respect to the beginning of the Reverse Data Packets field.

[00545] One embodiment of an example of advanced reverse data sampling is illustrated in graphical form in FIG. 52B. An ideal reverse data signal with zero round-trip delay is shown as a dotted-line waveform. The actual round-trip delay, between 3.5 and 4 MDDI_Stb cycles, can be observed as the difference in delay between solid waveform and the ideal. This is the same delay that would be measured using the Round Trip Delay Measurement Packet, and would be a measured round-trip delay value equal to 7 forward-link bit times. In this embodiment, reverse data bits are 2 MDDI_Stb pulses long, which is 4 forward-link bit times, which corresponds to a reverse rate divisor equal to 2. For advanced reverse data sampling it is convenient to use a pre-selected reverse rate divisor of 2 instead of computing it as described elsewhere. This appears to be a substantially optimum choice for advanced reverse data sampling because the ideal sampling point can easily be determined using the conventional measurements described above.

[00546] The ideal sampling point for reverse data can be easily computed by taking the remainder of the total round-trip delay divided by the number of forward link clocks per reverse bit, or round-trip delay modulo forward link clocks per reverse bit. Then subtract either 1 or 2 to get to a safe point away from the data transition. In this example, $7 \bmod 4 = 3$, then $3 - 1 = 2$, or $3 - 2 = 1$. The safe sampling point is either 1 or 2 forward link bit times from the edge of the "ideal" bit boundary for zero round-trip delay. The figure shows the sampling point at 2 forward link bit times from the ideal bit boundary, as indicated by the series of vertical arrows at the bottom of the timing diagram. The first sampling point is the first ideal bit boundary after the measured round-trip delay, plus the offset for safe sampling. In this example, the round trip delay measurement is 7, so the next ideal bit boundary is at the 8th bit time, then add either 1 or 2 for the safe sampling point, so the first bit shall be sampled at either 9 or 10 forward link bit times after the beginning of the Reverse Data Packets Field.

XI. Turn-Around and Guard Times

[00547] The Turn-Around 1 field in the Reverse Link Encapsulation Packet allows time for the host drivers to disable and the client drivers to enable simultaneously. The Guard Time 1 field in the Round Trip Delay Measurement Packet allows overlap of the host and client, so the client drivers can enable before the host interface drivers are

disabled. The Turn Around 2 field in the Reverse Link Encapsulation Packet allows data in the previous field from the client to be fully transmitted before the host drivers are enabled. The Guard Time 2 field provides a time value or period which allows the client and host drivers to drive simultaneously at a logic-zero level. The Guard Time 1 and Guard Time 2 fields are generally filled with pre-set or pre-selected values for lengths that are not meant to be adjusted. Depending on the interface hardware being used, these values may be developed using empirical data and adjusted in some instances to improve operation.

Turn-Around 1

- [00548] Several factors contribute to a determination of the length of Turn-Around 1 and these are the forward link data rate, the maximum disable time of the MDDI_Data drivers in the host, and the enable time of the client driver which is which is generally the same as the host disable time. The length of the Turn-Around 1 field is selected to be $24 \cdot t_{BTT}$. (Table XIII) The length in the number of forward link bytes of the Turn-Around 1 field is determined using the Interface Type Factor, and is computed using the relationship:

$$Length_{TurnAround1} = \frac{24}{8 \text{ bits / byte}} \cdot InterfaceTypeFactor_{FWD} = 3 \cdot InterfaceTypeFactor_{FWD}$$

where the Interface Type Factor is 1 for Type 1, 2 for Type 2, 4 for Type 3, and 8 for Type-4.

Turn-Around 2

- [00549] The factors that determine the length of time generally used for Turn Around 2 are, the round-trip delay of the communication link, the maximum disable time of the MDDI_Data drivers in the client, and the enable time of the host driver which is specified to be the same as the client driver disable time. The maximum host driver enable time and client driver disable time is specified elsewhere. The round-trip delay is measured in units of t_{BTT} . The minimum length specified in the number of forward link bytes of the Turn-Around 2 field is computed according to the relationship:

$$Length_{TurnAround2} \geq RoundUpToNextInteger \left(\frac{RoundTripDelay + 24}{8bits / byte} \cdot InterfaceTypeFactor_{FWD} \right)$$

[00550] For example, a Type 3 forward link with a round-trip delay of 10 forward link clocks typically uses a Turn Around 2 delay on the order of:

$$Length_{TurnAround2} \geq RoundUpToNextInteger \left(\frac{11 + 24}{8} \cdot 4 \right) = 18bytes$$

XII. Alternative Reverse Link Timing

[00551] While the use of timing and guard bands discussed above work to achieve a high data transfer rate interface, the inventors have discovered a technique to allow for reverse bit lengths that are shorter than the round trip time, by changing the reverse timing discovery.

[00552] As presented above, the previous approach to the timing of the reverse link is configured such that the number of clock cycles is counted from the last bit of the Guard Time 1 of a reverse timing packet until the first bit is sampled on the rising edge of an IO clock. That is the clock signal(s) used to time the inputs and outputs for the MDDI. The calculation for the reverse rate divisor is then given by:

$$reverse_rate_divisor = RoundUpToNextInteger \left(\frac{round_trip_delay + 1}{2} \right)$$

[00553] This provides a bit width equal to the round trip delay which results in a very reliable reverse link. However, the reverse link has been shown to be capable of running faster, or at a higher data transfer rate, which the inventors want to take advantage of. A new inventive technique allows utilizing additional capabilities of the interface to reach higher speeds.

[00554] This is accomplished by having the host count the number of clock cycles until a one is sampled, but with the host sampling the data line on both the rising and falling edges during the reverse timing packet. This allows the host to pick the most useful or even optimal sampling point within the reverse bit to ensure that the bit is stable. That

is, to find the most useful or optimal rising edge to sample data on for reverse traffic reverse encapsulation packets. The optimal sampling point depends on both the reverse link divisor and whether the first one was detected on a rising edge or a falling edge. The new timing method allows the host to just look for the first edge of the 0xFF 0xFF 0x00 pattern sent by the client for reverse link timing to determine where to sample in a reverse encapsulation packet.

[00555] Examples of the arriving reverse bit and how that bit would look for various reverse rate divisors, is illustrated in FIG. 64, along with a number of clock cycles that have occurred since the last bit of Guard Time 1. In Fig. 64, one can see that if the first edge occurs between a rising and falling edge (labeled as rise/fall), the optimal sampling point for a reverse rate divisor of one, the optimal sample point is a clock cycle edge labeled 'b', as that is the only rising edge occurring within the period of the reverse bit. For a reverse rate divisor of two, the optimal sampling point is probably still clock cycle leading edge 'b' as cycle edge 'c' is closer to a bit edge than 'b'. For a reverse rate divisor of four, the optimal sampling point is probably clock cycle edge 'd', as it is closer to the back edge of the reverse bit where the value has probably stabilized.

[00556] Returning to FIG. 64, if, however, the first edge occurs between a falling and rising edge (labeled as fall/rise), the optimal sampling point for a reverse rate divisor of one is sampling point clock cycle edge 'a', as that is the only rising edge within the reverse bit time period. For a reverse rate divisor of two, the optimal sampling point is edge 'b', and for a reverse rate divisor of four the optimal sampling point is edge 'c'.

[00557] One can see that as the reverse rate divisors get larger and larger, the optimal sampling point becomes easier to ascertain or select, as it should be the rising edge that is closest to the middle.

[00558] The host can use this technique to find the number of rising clock edges before the rising data edge of the timing packet data is observed on the data line. It can then decide, based on whether the edge occurs between a rising and falling edge or between a falling and rising edge, and what the reverse rate divisor is, how many additional clock cycles to add to a number counter, to reasonably ensure that the bit is always sampled as close to the middle as possible.

[00559] Once the host has selected or determined the number of clock cycles, it can "explore" various reverse rate divisors with the client to determine if a particular reverse rate divisor will work. The host (and client) can start with a divisor of one and check

the CRC of the reverse status packet received from the client to determine if this reverse rate functions appropriately to transfer data. If the CRC is corrupt, there is probably a sampling error, and the host can increase the reverse rate divisor and try to request a status packet again. If the second requested packet is corrupt, the divisor can be increased again and the request made again. If this packet is decoded correctly, this reverse rate divisor can be used for all future reverse packets.

[00560] This method is effective and useful because the reverse timing should not change from the initial round trip timing estimate. If the forward link is stable, the client should continue to decode forward link packets even if there are reverse link failures. Of course, it is still the responsibility of the host to set a reverse link divisor for the link, since this method does not guarantee a perfect reverse link. In addition, the divisor will depend primarily on the quality of the clock that is used to generate an IO clock. If that clock has a significant amount of jitter, there is a greater possibility of a sampling error. This error probability increases with the amount of clock cycles in the round trip delay.

[00561] This implementation appears to work best for Type 1 reverse data, but may present problems for Type 2 through Type 4 reverse data due to the skew between data lines potentially being too great to run the link at the rate that works best for just one data pair. However, the data rate probably does not need to be reduced to the previous method even with Type 2 through Type 4 for operation. This method may also work best if duplicated on each data line to select the ideal or an optimal clock sample location. If they are at the same sample time for each data pair, this method would continue to work. If they are at different sample periods, two different approaches may be used. The first is to select an desired or more optimized sample location for each data point, even if it is not the same for each data pair. The host can then reconstruct the data stream after sampling all of the bits from the set of data pairs: two bits for Type 2, four bits for Type 3, and eight bits for Type 4. The other option is for the host to increase the reverse rate divisor such that the data bits for every data pair can be sampled at the same clock edge.

XIII. Effects of Link Delay and Skew

[00562] Delay skew on the forward link between the MDDI_Data pairs and MDDI_Stb can limit the maximum possible data rate unless delay skew compensation is used. The differences in delay that cause timing skew are due to the controller logic, the line drivers and receivers, and the cable and connectors as outlined below.

A. Link Timing Analysis Limited by Skew (MDDI Type-1)

1. Delay and Skew Example of a Type 1 Link

[00563] A typical interface circuit similar to that shown in FIG. 41, is shown in FIG. 57 for accommodating a Type 1 interface link. In FIG. 57, exemplary or typical values for propagation delay and skew are shown for each of several processing or interface stages of an MDDI Type 1 forward link. Skew in the delay between MDDI_Stb and MDDI_Data0 causes the duty-cycle of the output clock to be distorted. Data at the D input of the receiver flip-flop (RXFF) stage using flip-flops 5728, 5732, changes slightly after the clock edge so that it can be sampled reliably. The figure shows two cascaded delay lines 5732a and 5732b being used to solve two different problems with creating this timing relationship. In the actual implementation these may be combined into a single delay element.

[00564] Data, Stb, and Clock Recovery Timing on a Type 1 Link for exemplary signal processing through the interface are illustrated in FIG. 58.

[00565] The total delay skew that is significant generally arises or comes from the sum of the skew in the following stages: transmitter flip-flop (TXFF) with flip-flops 5704, 5706; transmitter driver (TXDRVR) with drivers 5708, 5710; the CABLE 5702; receiver line receiver (RXRCVR) with receivers 5722, 5724; and receiver XOR logic (RXXOR). Delay1 5732a should match or exceed the delay of the XOR gate 5736 in the RXXOR stage which is determined by the relationship:

$$t_{PD - \min(\text{Delay } 1)} \geq t_{PD - \max(\text{XOR})}$$

[00566] It is desirable to meet this requirement so that the D input of receiver flip-flop 5728, 5732 does not change before its clock input. This is valid if the hold-time of RXFF is zero.

[005 67] The purpose or function of Delay2 is to compensate for the hold-time of the RXFF flip-flop according to the relationship:

$$t_{PD-\min(Delay2)} = t_{H(RXFF)}$$

[005 68] In many systems this will be zero because the hold time is zero, and of course in that case the maximum delay of Delay2 can also be zero.

[005 69] The worst-case contribution to skew in the receiver XOR stage is in the data-late/strobe-early case where Delay1 is at a maximum value and the clock output from the XOR gate comes as early as possible according to the relationship:

$$t_{SKEW-\max(RXOR)} = t_{PD-\max(Delay1)} - t_{PD-\min(XOR)}$$

[005 70] In this situation, the data may change between two bit periods, n and n+1, very close to the time where bit n+1 is clocked into the receiver flip-flop.

[005 71] The maximum data rate (minimum bit period) of an MDDI Type 1 link is a function of the maximum skew encountered through all the drivers, cable, and receivers in the MDDI link plus the total data setup into the RXFF stage. The total delay skew in the link up to the output of the RXRCVR stage can be expressed as:

$$t_{SKEW-\max(LINK)} = t_{SKEW-\max(TXFF)} + t_{SKEW-\max(TXDRVR)} + t_{SKEW-\max(CABLE)} + t_{SKEW-\max(RXRCVR)}$$

with the "cable" representing a variety of conductors or interconnections or wires and corresponding delay, and the minimum bit period is given by:

$$t_{BIT-\min} = t_{SKEW-\max(LINK)} + 2 \cdot t_{B-TP4} + t_{Asymmetry} + t_{SKEW-\max(RXOR)} + t_{filter-\text{hold}} + t_{PD-\max(Delay2)} + t_{SU(RXFF)}$$

[005 72] In the example shown in FIG. 57 for external mode, $t_{SKEW-\max(LINK)} = 1000 \text{ psec}$ and the minimum bit period can be expressed as:

$$t_{BIT-min} = 1000 + 2 \cdot 125 + 625 + 125 + 200 + 0 + 100 = 2300 \text{ psec},$$

or stated as approximately 434 Mbps. In the example shown in FIG. 57 for internal mode, $t_{SKEW-max(LINK)} = 500 \text{ psec}$ and the minimum bit period can be expressed as:

$$t_{BIT-min} = 500 + 2 \cdot 125 + 625 + 125 + 200 + 0 + 100 = 1800 \text{ psec},$$

or stated as approximately 555 Mbps.

B. Link Timing Analysis for MDDI Type 2, 3, and 4

[00573] A typical interface circuit similar to that shown in FIGs. 41 and 57, is shown in FIG. 59 for accommodating Type 2, 3, and 4 interface links. Additional elements are used in the TXFF (5904), TXDRVR (5908), RXRCVCR (5922), and RXFF (5932, 5928, 5930) stages to accommodate the additional signal processing. In FIG. 59, exemplary or typical values for propagation delay and skew are shown for each of several processing or interface stages of an MDDI Type 2 forward link. In addition to skew in the delay between MDDI_Stb and MDDI_Data0 affecting the duty-cycle of the output clock, there is also skew between both of these two signals and the other MDDI_Data signals. Data at the D input of the receiver flip-flop B (RXFFB) stage consisting of flip-flops 5928 and 5930, is changed slightly after the clock edge so it can be sampled reliably. If MDDI_Data1 arrives earlier than MDDI_Stb or MDDI_Data0 then MDDI_Data1 should be delayed to be sampled by at least the amount of the delay skew. To accomplish this, data is delayed using the Delay3 delay line. If MDDI_Data1 arrives later than MDDI_Stb and MDDI_Data0 and it is also delayed by Delay3 then the point where MDDI_Data1 changes is moved closer to the next clock edge. This process determines an upper limit of the data rate of an MDDI Type 2, 3, or 4 link. Some exemplary different possibilities for the timing or skew relationship of two data signals and MDDI_Stb with respect to each other is illustrated in FIGs. 60A, 60B, and 60C.

[00574] In order to sample data reliably in RXFFB when MDDI_DataX arrives as early as possible, Delay3 is set according to the relationship:

$$t_{PD-min(Delay3)} \geq t_{SKEW-max(LINK)} + t_{H(RXFFB)} + t_{PD-max(XOR)}$$

[00575] The maximum link speed is determined by the minimum allowable bit period. This is most affected when MDDI_DataX arrives as late as possible. In that case, the minimum allowable cycle time is given by:

$$t_{BIT-\min} = t_{SKEW-\max(LINK)} + t_{PD-\max(Delay3)} + t_{SU(RXFFB)} - t_{PD-\min(XOR)}$$

[00576] The upper bound of link speed is then:

$$t_{PD-\max(Delay3)} = t_{PD-\min(Delay3)}$$

and given that assumption:

$$t_{BIT-\min(lower-bound)} = 2 \cdot t_{SKEW-\max(LINK)} + t_{PD-\max(XOR)} + t_{SU(RXFFB)} + t_{H(RXFFB)}$$

[00577] In the example given above, the lower bound of the minimum bit period is given by the relationship:

$$t_{BIT-\min(lower-bound)} = 2 \cdot (1000 + 2 \cdot 125 + 625 + 200) + 1500 + 100 + 0 = 5750 \text{ psec,}$$

which is approximately 174 Mbps.

[00578] This is much slower than the maximum data rate that can be used with a Type 1 link. The automatic delay skew compensation capability of MDDI significantly reduces the affect that delay skew has on the maximum link rate factor is just on-the-edge of valid data setup. The calibrated skew between MDDI_Data0 and MDDI_Stb is:

$$t_{SKEW-\max(Calibrated)} = 2 \cdot t_{TAP-SPACING-\max},$$

and the minimum bit period is:

$$t_{BIT-\min-Calibrated} = t_{SKEW-\max(Calibrated)} + 2 \cdot t_{B-TP4} + t_{Asymmetry} + t_{jitter-host} + t_{SKEW-\max(RXAND+RXOR)} + t_{SU(RXFFB)}$$

[00579] Where “TB” or t_B represents signal jitter from a bit boundary to minimum output level. Asymmetry simply refers to the asymmetrical nature of internal delay through or of the differential receivers. “TP4” is associated with or is effectively defined for electrical characterization and testing purposes as the connection or interface (pins of the MDDI controller device in the client) for the differential line drivers and receivers for the client. It represents a convenient or predetermined point from which signal delay is measured and characterized for the link throughout the rest of a system. In one embodiment, a maximum value of the parameter t_B at TP4 is defined by the relationship

$$t_{\text{Differential-Skew-TP4-DRVR-EXT}} = 0.3 \cdot t_{\text{BIT}} \quad \text{for the external mode and}$$

$$t_{\text{Differential-Skew-TP4-DRVR-INT}} = 0.6 \cdot t_{\text{BIT}} \quad \text{for the internal mode for the client transmitters; and}$$

$$t_{\text{B-TP4-RCVR-EXT}} = 0.051 \cdot t_{\text{BIT}} + 175 \text{ ps} \quad \text{for the external mode for the client receivers.}$$

[00580] The label TP4 being simply useful in numbering various test points (TP) in the interface and links. In one embodiment, this test point location is defined to be the same for both internal and external modes. There is a corresponding “TP0” test point for, or associated with, the connection or interface pins of the MDDI controller device in the host that contains the differential line drivers and receivers. In this embodiment, a maximum value of the parameter T_B at TP0 is defined by the relationship

$$t_{\text{B-TP0-RCVR-INT}} = 0.051 \cdot t_{\text{BIT}} + 50 \text{ ps}, \quad \text{for the internal mode, and}$$

$$t_{\text{B-TP0-RCVR-EXT}} = 0.051 \cdot t_{\text{BIT}} + 175 \text{ ps} \quad \text{for the external mode for the host receivers; and}$$

$$t_{\text{B-TP0}} = 0.102 \cdot t_{\text{BIT}} \quad \text{for the host transmitters.}$$

[00581] In the example shown in FIG. 59, $t_{\text{SKEW-max(Data0-Stb-Calibrated)}} = 300 \text{ psec}$ and the minimum bit period:

$$t_{\text{BIT-min-Calibrated}} = 300 + 2 \cdot 125 + 625 + 200 + 175 + 100 = 1650 \text{ psec},$$

approximately 606 Mbps.

[00582] In order to sample data reliably in RXFFB when MDDI_Data1 arrives as early as possible, the associated programmable delay is adjusted to the optimal setting with an accuracy of one tap, and an additional tap delay is added for safety. The maximum link speed is determined by the minimum allowable bit period. This is most affected when

MDDI_Data1 arrives as late as possible. In that case the minimum allowable cycle time is:

$$t_{BIT-min-Data1-Calibrated} = 2 \cdot t_{TAP-Spacing-max} + 2 \cdot t_{TA-TP4},$$

where "TA" or t_A represents signal jitter from a bit boundary to center crossing.

[00583] In the example given in FIG. 59, the lower bound of the minimum bit period based on sampling MDDI_Data1 is:

$$t_{BIT-min-Data1-Calibrated} = 2 \cdot 150 + 2 \cdot 125 = 550 \text{ p sec}$$

[00584] In one embodiment, a typical total delay time for delay skew, delay asymmetry, and Clock Jitter in the host transmitter for Internal Mode would be defined as:

$$t_{Asymmetry-TXFF} + t_{Asymmetry-TXDRVR} + t_{Skew-TXFF} + t_{Skew-TXDRVR} + t_{jitter-host} = 0.467 \cdot (t_{BIT} - 150 \text{ ps}),$$

and for the external mode as:

$$t_{Asymmetry-TXFF} + t_{Asymmetry-TXDRVR} + t_{Skew-TXFF} + t_{Skew-TXDRVR} + t_{jitter-host} = 0.1BD \cdot (t_{BIT} - 150 \text{ TBDps})$$

while a typical total delay time for delay skew, delay asymmetry, and setup time in the client device (t_{B-TP4}) for internal mode is:

$$t_{Asymmetry-RXRCVR} + t_{Asymmetry-RXXOR} + t_{Skew-RXRCVR} + t_{Skew-RXXOR} + t_{setup-RXFF} = 0.307 \cdot (t_{BIT} - 150 \text{ ps}),$$

and for the external mode:

$$t_{Asymmetry-RXRCVR} + t_{Asymmetry-RXXOR} + t_{Skew-RXRCVR} + t_{Skew-RXXOR} + t_{setup-RXFF} = 0.1BD \cdot (t_{BIT} - \text{TBDps}),$$

where the term TBD is a flexible place keeping label for future to be determined values which will depend on a variety of well understood characteristics and operational requirements for the external mode connections.

XIV. Physical Layer Interconnection Description

[00585] Physical connections useful for implementing an interface according to the present invention can be realized using commercially available parts such as part number 3260-8S2(01) as manufactured by Hirose Electric Company Ltd. on the host side, and part number 3240-8P-C as manufactured by Hirose Electric Company Ltd. on the client device side. An exemplary interface pin assignment or "pinout" for such connectors used with a Type-1/Type 2 interfaces is listed in Table XV, and illustrated in FIG. 61.

Table XV

Signal Name	Pin Number	Signal Name	Pin Number
MDDI_Pwr	1	MDDI_Gnd	11
MDDI_Stb+	2	MDDI_Stb-	12
MDDI_Data0+	4	MDDI_Data0-	14
MDDI_Data1+	6	MDDI_Data1-	16
MDDI_Data2+	8	MDDI_Data2-	18
MDDI_Data3+	10	MDDI_Data3-	20
MDDI_Data4+	9	MDDI_Data4-	19
MDDI_Data5+	7	MDDI_Data5-	17
MDDI_Data6+	5	MDDI_Data6-	15
MDDI_Data7+	3	MDDI_Data7-	13
		Shield	

[00586] The shield is connected to the HOST_Gnd in the host interface, and a shield drain wire in the cable is connected to the shield of the client connector. However, the shield and drain wire are not connected to the circuit ground inside of a client.

[00587] Interconnection elements or devices are chosen or designed in order to be small enough for use with mobile communication and computing devices, such as PDAs and wireless telephones, or portable game devices, without being obtrusive or unaesthetic in comparison to relative device size. Any connectors and cabling should be durable enough for use in the typical consumer environment and allow for small size, especially for the cabling, and relatively low cost. The transfer elements should accommodate data and strobe signals that are differential NRZ data having a transfer rate up to around 450 Mbps for Type 1 and Type 2 and up to 3.6 Gbps for the 8-bit parallel Type 4 version.

[00588] For internal mode applications there are either no connectors in the same sense for the conductors being used or such connection elements tend to be very miniaturized. One example is zero insertion force “sockets” for receiving integrated circuits or elements housing either the host or client device. Another example is where the host and client reside on printed circuit boards with various interconnecting conductors, and have “pins” or contacts extending from housings which are soldered to contacts on the conductors for interconnection of integrated circuits.

XV. Operation

[00589] A summary of the general steps undertaken in processing data and packets during operation of an interface using embodiments of the invention is shown in FIGS. 54A and 54B, along with an overview of the interface apparatus processing the packets in FIG. 55. In these figures, the process starts in a step 5402 with a determination as to whether or not the client and host are connected using a communication path, here a cable. This can occur through the use of periodic polling by the host, using software or hardware that detects the presence of connectors or cables or signals at the inputs to the host (such as is seen for USB interfaces), or other known techniques. If there is no client connected to the host, then it can simply enter a wait state of some predetermined length, depending upon the application, go into a hibernation mode, or be inactivated to await future use which might require a user to take action to reactivate the host. For example, when a host resides on a computer type device, a user might have to click on a screen icon or request a program that activates the host processing to look for the client. Again, simple plug in of a USB type connection could activate host processing, depending on the capabilities and configuration of the host or resident host software.

[00590] Once a client is connected to the host, or visa versa, or detected as being present, either the client or the host sends appropriate packets requesting service in steps 5404 and 5406. The client could send either Client Service Request or Status packets in step 5404. It is noted that the link, as discussed above, could have been previously shut down or be in hibernation mode so this may not be a complete initialization of the communication link that follows. Once the communication link is synchronized and the host is trying to communicate with the client, the client also provides a Client

Capabilities packet to the host, as in step 5408. The host can now begin to determine the type of support, including transfer rates, the client can accommodate.

[00591] Generally, the host and client also negotiate the type (rate/speed) of service mode to be used, for example Type 1, Type 2, and so forth, in a step 5410. Once the service type is established the host can begin to transfer information. In addition, the host may use Round Trip Delay Measurement Packets to optimize the timing of the communication links in parallel with other signal processing, as shown in step 5411.

[00592] As stated earlier, all transfers begin with a Sub-Frame Header Packet, shown being transferred in step 5412, followed by the type of data, here video and audio stream packets, and filler packets, shown being transferred in step 5414. The audio and video data will have been previously prepared or mapped into packets, and filler packets are inserted as needed or desired to fill out a required number of bits for the media frames. The host can send packets such as the Forward Audio Channel Enable Packets to activate sound devices. In addition, the host can transfer commands and information using other packet types discussed above, here shown as the transfer of Color Map, Bit Block Transfer or other packets in step 5416. Furthermore, the host and client can exchange data relating to a keyboard or pointing devices using the appropriate packets.

[00593] During operation, one of several different events can occur which lead to the host or client desiring a different data rate or type of interface mode. For example, a computer or other device communicating data could encounter loading conditions in processing data that causes a slow down in the preparation or presentation of packets. A client device receiving the data could change from a dedicated AC power source to a more limited battery power source, and either not be able to transfer in data as quickly, process commands as readily, or not be able to use the same degree of resolution or color depth under the more limited power settings. Alternatively, a restrictive condition could be abated or disappear allowing either device to transfer data at higher rates. This being more desirable, a request can be made to change to a higher transfer rate mode.

[00594] If these or other types of known conditions occur or change, either the host or client may detect them and try to renegotiate the interface mode. This is shown in step 5420, where the host sends Interface Type Handoff Request Packets to the client requesting a handoff to another mode, the client sends Interface Type Acknowledge Packets confirming a change is sought, and the host sends Perform Type Handoff Packets to make the change to the specified mode.

- [00595] Although, not requiring a particular order of processing, the client and host can also exchange packets relating to data intended for or received from pointing devices, keyboards, or other user type input devices associated primarily with the client, although such elements may also be present on the host side. These packets are typically processed using a general processor type element and not the state machine (5502). In addition, some of the commands discussed above will also be processed by the general processor. (5504, 5508)
- [00596] After data and commands have been exchanged between the host and client, at some point a decision is made as to whether or not additional data is to be transferred or the host or client is going to cease servicing the transfer. This is shown in step 5422. If the link is to enter either a hibernation state or be shut down completely, the host sends a Link Shutdown packet to the client, and both sides terminate the transfer of data.
- [00597] The packets being transferred in the above operations processing will be transferred using the drivers and receivers previously discussed in relation to the host and client controllers. These line drivers and other logic elements are connected to the state machine and general processors discussed above, as illustrated in the overview of FIG. 55. In Fig. 55, a state machine 5502 and general processors 5504 and 5508 may further be connected to other elements not shown such as a dedicated USB interface, memory elements, or other components residing outside of the link controller with which they interact, including, but not limited to, the data source, and video control chips for view display devices.
- [00598] The processors, and state machine provide control over the enabling and disabling of the drivers as discussed above in relation to guard times, and so forth, to assure efficient establishment or termination of communication link, and transfer of packets.

XVI. Display Frame Buffers

- [00599] Video data buffering requirements are different for moving video images compared to computer graphics. Pixel data is most often stored in a local frame buffer in the client so the image on the client can be refreshed locally.
- [00600] When full-motion video is being displayed (nearly every pixel in the display changes each Media Frame) it is usually preferred to store the incoming pixel data in

one frame buffer while the image on the display is being refreshed from a second frame buffer. More than two display buffers may be used to eliminate visible artifacts as described below. When an entire image has been received in one frame buffer then the roles of the buffers can be swapped, and the newly received image is then used to refresh the display and the other buffer is filled with the next frame of the image. This concept is illustrated in FIG. 88A, where pixel data is written to the offline image buffer by setting the Display Update bits to "01."

[00601] In other applications the host needs to update only a small portion of the image without having to repaint the entire image. In this situation it is desired to write the new pixels directly to the buffer being used to refresh the display, as illustrated in detail FIG. 88B.

[00602] In applications that have a fixed image with a small video window it is easiest to write the fixed image to both buffers (display update bits equal to "11") as shown in FIG. 88C, and subsequently write the pixels of the moving image to the offline buffer by setting the display update bits to "01."

[00603] The following rules describe the useful manipulation of buffer pointers while simultaneously writing new information to the client and refreshing the display. Three buffer pointers exist: `current_fill` points to the buffer currently being filled from data over the MDDI link. `Just_filled` points to the buffer that was most recently filled. `being_displayed` points to the buffer currently being used to refresh the display. All three buffer pointers may contain values from 0 to N-1 where N is the number of display buffers, and $N \geq 2$. Arithmetic on buffer pointers is $\text{rmod } N$, e.g. when $N=3$ and `current_fill=2`, incrementing `current_fill` causes `current_fill` to be set to 0. In the simple case where $N=2$, `just_filled` is always the complement of `current_fill`. On every MDDI Media Frame boundary (Sub-frame Header Packet with the Sub-frame Count field equal so zero) perform the following operations in the order specified: set `just_filled` equal to `current_fill`, and set `current_fill` equal to `current_fill + 1`.

[00604] MDDI Video Stream Packets update the buffers according to the structure or methodology of: when Display Update Bits equal to '01', pixel data is written to the buffer specified by `current_fill`; when Display Update Bits equal to '00', pixel data is written to the buffer specified by `just_filled`; and when Display Update Bits equal to "11," pixel data is written to all buffers. The display is refreshed from the buffer specified by the `being_displayed` pointer. After the display refreshes the last pixel in

one frame refresh epoch and before it begins to refresh the first pixel in the next frame refresh epoch the display update process performs the operation of setting being_refreshed equal to just_filled.

[00605] The Packets with a Pixel Data Attribute field contain a pair of Display Update Bits that specify the frame buffer where the pixel data is to be written. The Client Capability Packet has three additional bits that indicate which combinations of the Display Update Bits are supported in the client. In many cases, computer-generated images need to be incrementally updated based on user input or derived from information received from a computer network. Display Update Bit combinations "00" and "11" support this mode of operation by causing the pixel data to be written to the frame buffer being displayed or to both frame buffers.

[00606] When accommodating video images, FIG. 89 illustrates how video images are displayed using a pair of frame buffers when video data is transmitted over the MDDI link with the Display Update Bits equal to "01." After a media-frame boundary is detected on the MDDI link, the display refresh process will begin refreshing from the next frame buffer when the refresh process for the frame currently being refreshed is completed.

[00607] An important assumption related to FIG. 89 is that the image is received from the host as a continuous stream of pixels that are transmitted in the same order that the client uses to read the pixels from the frame buffer to refresh the display (usually upper-left, reading row by row, to the bottom-right corner of the screen. This is an important detail in the cases where the Display Refresh and Image Transfer operations reference the same frame buffer.

[00608] It is necessary for the display refresh frame rate to be greater than the image transfer frame rate to avoid displaying partial images. FIG. 90 shows how image fragmentation can occur with a slow display refresh rate that is the display refresh is slower than the image transfer.

[00609] In an image that contains a combination of computer graphic images and moving video pictures the video pixel data might occupy a small portion of a media-frame. This could be significant in situations where the display refresh operation and the image transfer reference the same frame buffer. These situations are shown by a cross-hatched shading in FIG. 91, where the pixels read from the buffer to refresh the display might be

the pixels written to the buffer two frames ago, or they may correspond to the frame immediately being written to the same frame buffer.

[00610] The use of three frame buffers in the client will resolve the problem of the small window of contention for access to a frame buffer as shown in FIG. 92.

[00611] However, there is still a problem if the display refresh rate is less than the media-frame rate over the MDDI link as shown in FIG. 93.

[00612] The use of a single buffer for moving video images is somewhat problematic as shown FIG. 94. With the display refresh faster than the image transfer into the buffer, the image being refreshed sometimes will show the upper portion of the frame being written and the lower portion of the image will be the frame previously transferred. With the display refresh faster than the image transfer (the preferred mode of operation) there will be more frequent instances of frames showing a similar split image.

XVII. Multiple Client Support

[00613] The current protocol version does not appear to directly support multiple client devices. However, most packets contain a reserved Client ID field that can be used to address specific client devices in a system with multiple clients. Currently, for many applications this client ID or these client IDs are set to be zero. The sub-frame header packet also contains a field to indicate whether or not the host supports a multiple client system. Therefore, there is a manner in which multiple client devices would likely be connected and addressed in future applications of the MDDI or protocol to aid system designers to plan for future compatibility with multiple client hosts and clients.

[00614] In systems having multiple clients it is useful for clients to be connected to the host using a daisy-chain of clients, or using hubs, as shown in Fig. 95, or using a combination of these techniques as shown in FIG. 96. It also may be useful for a host to display certain error messages to manage the connected clients, such as an error message when one or more clients desiring address 0 are connected, which should not be the case for multi-client systems, as such displays expect to be or are set to operate as the only client.

XIII. Addendum

[00615] In addition to the formats, structures, and contents discussed above for the various packets used to implement the architecture and protocol for embodiments of the invention, more detailed field contents or operations are presented here for some of the packet types. These are presented here to further clarify their respective use or operations to enable those skilled in the art to more readily understand and make use of the invention for a variety of applications. Only a few of the fields not already discussed are discussed further here. In addition, these fields are presented with exemplary definitions and values in relation to the embodiments presented above. However, such values are not to be taken as limitations of the invention, but represent one or more embodiments useful for implementing the interface and protocol, and not all embodiments need be practiced together or at the same time. Other values can be used in other embodiments to achieve the desired presentation of data or data rate transfer results, as will be understood by those skilled in the art.

A. For Video Stream Packets

[00616] In one embodiment, the Pixel Data Attributes field (2 byte) has a series of bit values that are interpreted as follows. Bits 1 and 0 select how the display pixel data is routed. For bit values of '11' pixel data is displayed to or for both eyes, for bit values '10', pixel data is routed only to the left eye, and for bit values '01', pixel data is routed only to the right eye, and for bit values of '00' the pixel data is routed to an alternate display as may be specified by bits 8 through 11 discussed below. If the primary display in or being used or operated by a client does not support stereo images or imaging in some form, then these commands cannot effectively be implanted to have an impact as desired by the display. In this situation or configuration the client should route pixel data to a primary display regardless of the bit values or for any of the bit combinations '01,' '10,' or '11,' since the resulting commands or control won't be implemented by the display. It is recommended, but not required by the embodiments, that the value '11' be used to address the primary display in those clients that do not support stereo display capability.

[00617] Bit 2 indicates whether or not the Pixel Data is presented in an interlace format, with a value of '0' meaning the pixel data is in the standard progressive format, and that the row number (pixel Y coordinate) is incremented by 1 when advancing from one row to the next. When this bit has a value of '1', the pixel data is in interlace format, and the row number is incremented by 2 when advancing from one row to the next. Bit 3 indicates that the Pixel Data is in alternate pixel format. This is similar to the standard interlace mode enabled by bit 2, but the interlacing is vertical instead of horizontal. When Bit 3 is '0' the Pixel Data is in the standard progressive format, and the column number (pixel X coordinate) is incremented by 1 as each successive pixel is received. When Bit 3 is '1' the Pixel Data is in alternate pixel format, and the column number is incremented by 2 as each pixel is received.

[00618] Bit 4 indicates whether or not the Pixel data is related to a display or a camera, as where data is being transferred to or from an internal display for a wireless phone or similar device or even a portable computer, or such other devices as discussed above, or the data is being transferred to or from a camera built into or directly coupled to the device. When Bit 4 is '0' the Pixel data is being transferred to or from a display frame buffer. When Bit 4 is '1' Pixel data is being transferred to or from a camera or video device of some type, such devices being well known in the art.

- [00619] Bit 5 is used to indicate when the pixel data contains the next consecutive row of pixels in the display. This is considered the case when Bit 5 is set equal to '1'. When bit 5 is set to '1' then the X Left Edge, Y Top Edge, X Right Edge, Y Bottom Edge, X Start, and Y Start parameters are not defined and are ignored by the client. When Bit 5 is set at a logic-one level, this indicates that the pixel data in this packet is the last row of pixels in the image. Bit 8 of the Client Feature Capability Indicators field of the Client Capability Packet indicates whether this feature is supported.
- [00620] Bits 7 and 6 are Display Update Bits that specify a frame buffer where the pixel data is to be written. The more specific effects are discussed elsewhere. For bit values of '01' Pixel data is written to the offline image buffer. For bit values of '00' Pixel data is written to the image buffer used to refresh the display. For bit values of '11' Pixel data is written to all image buffers. The bit values or combination of '10' is treated as an invalid value or designation and Pixel data is ignored and not written to any of the image buffers. This value may have use for future applications of the interface.
- [00621] Bits 8 through 11 form a 4-bit unsigned integer that specifies an alternate display or display location where pixel data is to be routed. Bits 0 and 1 are set equal to '00' in order for the display client to interpret bits 8 through 11 as an alternate display number. If bits 0 and 1 are not equal to '00' then bits 8 through 11 are set to logic-zero levels.
- [00622] Bits 12 through 14 are reserved for future use and are generally set to logic-zero levels. Bit 15, as discussed, is used in conjunction with bit 5, and setting bit 15 to logic-one indicates that the row of pixels in the Pixel Data field is the last row of pixels in a frame of data. The next Video Stream Packet having bit 5 set to logic-one will correspond to the first row of pixels of the next video frame.
- [00623] The 2-byte X Start and Y Start fields specify the absolute X and Y coordinates of the point (X Start, Y Start) for the first pixel in the Pixel Data field. The 2-byte X Left Edge and Y Top Edge fields specify the X coordinate of the left edge and Y coordinate of the top edge of the screen window filled by the Pixel Data field, while the X Right Edge and Y Bottom Edge fields specify the X coordinate of the right edge, and the Y coordinate of the bottom edge of the window being updated.
- [00624] The Pixel Count field (2 bytes) specifies the number of pixels in the Pixel Data field below.
- [00625] The Parameter CRC field (2 bytes) contains a CRC of all bytes from the Packet Length to the Pixel Count. If this CRC fails to check then the entire packet is discarded.

- [00626] The Pixel Data field contains the raw video information that is to be displayed, and which is formatted in the manner described by the Video Data Format Descriptor field. The data is transmitted one "row" at a time as discussed elsewhere. When Bit 5 of the Pixel Data Attributes field is set at logic level one, then the Pixel Data field contains exactly one row of pixels, with the first pixel being transmitted corresponding to the left-most pixel and the last pixel transmitted corresponding to the right-most pixel.
- [00627] The Pixel Data CRC field (2 bytes) contains a 16-bit CRC of only the Pixel Data. If a CRC verification of this value fails then the Pixel Data can still be used, but the CRC error count is incremented.

B. For Audio Stream Packets

- [00628] In one embodiment, the Audio Channel ID field (1 byte) uses an 8 bit unsigned integer value to identify a particular audio channel to which audio data is sent by the client device. The physical audio channels are specified in or mapped to physical channels by this field as values of 0, 1, 2, 3, 4, 5, 6, or 7 which indicate the left front, right front, left rear, right rear, front center, sub-woofer, surround left, and surround right channels, respectively. An audio channel ID value of 254 indicates that the single stream of digital audio samples is sent to both the left front and right front channels. This simplifies communications for applications such as where a stereo headset is used for voice communication, productivity enhancement apps are used on a PDA, or other applications where a simple User Interface generates warning tones. Values for the ID field ranging from 8 through 253, and 255 are currently reserved for use where new designs desire additional designations, as anticipated by those skilled in the art.
- [00629] The Reserved 1 field (1 byte) is generally reserved for future use, and has all bits in this field set to zero. One function of this field is to cause all subsequent 2 byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.
- [00630] The Audio Sample Count field (2 bytes) specifies the number of audio samples in this packet.
- [00631] The Bits Per Sample and Packing field contains 1 byte that specifies the packing format of audio data. In one embodiment, the format generally employed is for Bits 4 through 0 to define the number of bits per PCM audio sample. Bit 5 then specifies

whether or not the Digital Audio Data samples are packed. As mentioned above, FIG. 12 illustrates the difference between packed and byte-aligned audio samples. A value of '0' for Bit 5 indicates that each PCM audio sample in the Digital Audio Data field is byte-aligned with the interface byte boundary, and a value of '1' indicates that each successive PCM audio sample is packed up against the previous audio sample. This bit is effective only when the value defined in bits 4 through 0 (the number of bits per PCM audio sample) is not a multiple of eight. Bits 7 through 6 are reserved for use where system designs desire additional designations and are generally set at a value of zero.

[00632] The Audio Sample Rate field (1 byte) specifies the audio PCM sample rate. The format employed is for a value of 0 to indicate a rate of 8,000 samples per second (sps), a value of 1 indicates 16,000 sps., value of 2 for 24,000 sps, value of 3 for 32,000 sps, value of 4 for 40,000 sps, value of 5 for 48,000 sps, value of 6 for 11,025 sps, value of 7 for 22,050 sps, and value of 8 for 44,100 sps, respectively, with values of 9 through 255 being reserved for future use, so they are currently set to zero.

[00633] The Parameter CRC field (2 bytes) contains a 16-bit CRC of all bytes from the Packet Length to the Audio Sample Rate. If this CRC fails to check appropriately, then the entire packet is discarded. The Digital Audio Data field contains the raw audio samples to be played, and is usually in the form of a linear format as unsigned integers. The Audio Data CRC field (2 bytes) contains a 16-bit CRC of only the Audio Data. If this CRC fails to check, then the Audio Data can still be used, but the CRC error count is incremented.

C. For User-Defined Stream Packets

[00634] In one embodiment, the 2-byte Stream ID Number field is used to identify a particular user defined stream. The contents of the Stream Parameters and Stream Data fields, are typically defined by the MDDI equipment manufacturer. The 2-byte Stream Parameter CRC field contains a 16-bit CRC of all bytes of the stream parameters starting from the Packet Length to the Audio Coding byte. If this CRC fails to check, then the entire packet is discarded. Both the Stream Parameters and Stream Parameter CRC fields may be discarded if not needed by an end application of the MDDI, that is, they are considered optional. The 2-byte Stream Data CRC field contains a CRC of only the Stream Data. If this CRC fails to check appropriately, then use of the Stream

Data is optional, depending on the requirements of the application. Use of the stream data contingent on the CRC being good, generally requires that the stream data be buffered until the CRC is confirmed as being good. The CRC error count is incremented if the CRC does not check.

D. For Color Map Packets

[00635] The 2-byte hClient ID field contains information or values that are reserved for a Client ID, as used previously. Since this field is generally reserved for future use, the current value is set to zero, by setting the bits to '0'.

[00636] The 2-byte Color Map Item Count field uses values to specify the total number of 3-byte color map items that are contained in the Color Map Data field, or the color map table entries that exist in the Color Map Data in this packet. In this embodiment, the number of bytes in the Color Map Data is 3 times the Color Map Item Count. The Color Map Item Count is set equal to zero to send no color map data. If the Color Map Size is zero then a Color Map Offset value is generally still sent but it is ignored by the display. The Color Map Offset field (4 bytes) specifies the offset of the Color Map Data in this packet from the beginning of the color map table in the client device.

[00637] A 2-byte Parameter CRC field contains a CRC of all bytes from the Packet Length to the Audio Coding byte. If this CRC fails to check then the entire packet is discarded.

[00638] For the Color Map Data field, the width of each color map location is a specified by the Color Map Item Size field, where in one embodiment the first part specifies the magnitude of blue, the second part specifies the magnitude of green, and the third part specifies the magnitude of red. The Color Map Size field specifies the number of 3-byte color map table items that exist in the Color Map Data field. If a single color map cannot fit into one Video Data Format and Color Map Packet, then the entire color map may be specified by sending multiple packets with different Color Map Data and Color Map Offsets in each packet. The number of bits of blue, green, and red in each color map data item is generally the same as specified in the Color Map RGB Width field of the Display Capability Packet.

[00639] A 2-byte Color Map Data CRC field contains a CRC of only the Color Map Data. If this CRC fails to check then the Color Map Data can still be used but the CRC error count is incremented.

- [00640] Each color map data item is to be transmitted in the order: blue, green, red, with the least significant bit of each component transmitted first. The individual red, green, and blue components of each color map item are packed, but each color map item (the least significant bit of the blue component) should be byte-aligned. Fig. 97 illustrates an example of color map data items with 6 bits of blue, 8 bits of green, and 7 bits of red. For this example, the Color Map Item Size in the Color Map Packet is equal to 21, and the Color Map RGB Width field of the Client Capability Packet is equal to 0x0786.

E. For Reverse Link Encapsulation Packets

- [00641] The Parameter CRC field (2 bytes) contains a 16-bit CRC of all bytes from the Packet Length to the Turn-Around Length. If this CRC fails to check, then the entire packet is discarded.
- [00642] In one embodiment, the Reverse Link Flags field (1 byte) contains a set of flags to request information from the client and specify a reverse link type. If a bit (for example, Bit 0) is set to a logic-one level, then the host requests the specified information from the client, but if the bit is set to a logic-zero level then the host does not need the information from the client. Bit 0 is used to indicate when the host desires the Client Capability Packet, which is generally sent by the client to the host in the Reverse Data Packets field. Bit 1 is used to indicate when the host desires the Client Request and Status Packet which is sent by the client to the host in the Reverse Data Packets field. The remaining bits (here Bits 2 through 7) are reserved for future use and are set to zero. However, more bits can be used as desired to set flags for the reverse link.
- [00643] The Reverse Rate Divisor field (1 byte) specifies the number of MDDI_Stb cycles that occur in relation to the reverse link data clock. The reverse link data clock is equal to the forward link data clock divided by two times the Reverse Rate Divisor. The reverse link data rate is related to the reverse link data clock and the Interface Type on the reverse link. In this embodiment, for a Type 1 interface the reverse data rate equals the reverse link data clock, for Type 2, Type 3, and Type 4 interfaces the reverse data rates equal two times, four times, and eight times the reverse link data clock, respectively.
- [00644] The All Zero 1 field contains a group of bytes, here 8, that is set equal to zero in value by setting the bits at a logic-zero level, and is used to ensure that all MDDI_Data

signals are at a logic-zero level for a sufficient time to allow the client to begin recovering clock using only MDDI_Stb prior to disabling the host's line drivers during the Turn-Around 1 field. In one embodiment, the length of the All Zero 1 field is greater than or equal to the number of forward link byte transmission times in the round-trip delay of the cable.

[00645] The Turn-Around 1 Length field (1 byte) specifies the total number of bytes that are allocated for Turn-Around 1, establishing the first turn-around period. The Turn-Around 1 field employs the number of bytes specified by the Turn-Around 1 Length parameter are allocated to allow the MDDI_Data line drivers in the client to enable, before the line drivers in the host are disabled. The client enables its MDDI_Data line drivers during bit 0 of Turn-Around 1 and the host disables its outputs so as to be completely disabled prior to the last bit of Turn-Around 1. The MDDI_Stb signal behaves as though MDDI_Data0 were at a logic-zero level during the entire Turn Around 1 period. A more complete description of the setting of Turn-Around 1 is given above.

[00646] The Reverse Data Packets field contains a series of data packets transferred from the client to host. The client may send filler packets or drive the MDDI_Data lines to a logic-zero state or level when it has no data to send to the host. In this embodiment, if the MDDI_Data lines are driven to zero, the host will interpret this as a packet with a zero length (not a valid length) and the host will accept no additional packets from the client for the duration of the current Reverse Link Encapsulation Packet.

[00647] The Turn-Around 2 Length field (1 byte) specifies the total number of bytes that are allocated for Turn-Around 2, for establishing a second turn-around period. The recommended length of Turn-Around 2 is the number of bytes required for the round-trip delay plus the time required for the host to enable its MDDI_Data drivers. Turn-Around 2 Length may be also be a value larger than the minimum required or calculated value to allow sufficient time to process reverse link packets in the host.

[00648] The Turn Around 2 field consists of the number of bytes as specified by the Turn-Around Length parameter. The host waits for at least the round trip delay time before it enables its MDDI_Data line drivers during Turn-Around 2. The host enables its MDDI_Data line drivers so that they are generally completely enabled prior to the last bit of Turn-Around 2, and the client disables its outputs so that they are generally completely disabled prior to the last bit of Turn-Around 2. The purpose of the Turn-

Around 2 field is to allow the remaining amount of data from the Reverse Data Packets field to be transmitted or transferred from the client. Variations in different systems implementing the interface and the amount of safety margin allocated, it is possible that neither the host nor client will be driving the MDDI_Data signals to a logic-zero level during some parts of the Turn Around 2 field period, as seen by the line receivers in or at the host. The MDDI_Stb signal behaves as though the MDDI_Data0 were at a logic-zero level during substantially the entire Turn Around 2 period. A description of the setting of Turn-Around 2 is given above.

[00649] The Reverse Data Packets field contains a series of data packets being transferred from the client to a host. As stated earlier, Filler packets are sent to fill the remaining space that is not used by other packet types.

[00650] The All Zero 2 field contains a group of bytes (8 in this embodiment) that is set equal to zero in value by setting the bits at a logic-zero level, and is used to ensure that all MDDI_Data signals are at a logic-zero level for a sufficient time to allow the client to begin recovering clock using both MDDI_Data0 and MDDI_Stb after enabling the host's line drivers following the Turn-Around 2 field.

F. For Client Capability Packets

[00651] As illustrated for one embodiment, the Protocol Version field uses 2 bytes to specify a protocol version used by the client. The initial version is currently set equal to one, and will be changed over time as new versions are generated as would be known, while the Minimum Protocol Version field uses 2 bytes to specify the minimum protocol version that the client can employ or interpret. In this case, a zero value is also a valid value. The Data Rate Capability field (2 bytes) specifies the maximum data rate the client can receive on each data pair on the forward link of the interface, and is specified in the form of megabits per second (Mbps). The Interface Type Capability field (1 byte) specifies the interface types that are supported on the forward and reverse links. A bit set to '1' indicates that a specified interface type is supported, and a bit set to '0' indicates that the specified type is not supported. Hosts and clients should support at least Type 1 on the forward and reverse links. There is no requirement to support a contiguous range of interface types. For example, it would be perfectly valid to support only Type 1 and Type 3, but not Type 3 and Type 4 in an interface. It is also not necessary for the forward and reverse links to operate with the same interface type.

However, when a link comes out of hibernation both forward and reverse links should commence operating in Type 1 mode, until other modes may be negotiated, selected, or otherwise approved for use by both the host and client.

[00652] The supported interfaces are indicated in one embodiment by selecting Bit 0, Bit 1, or Bit 2 to select either a Type 2 (2 bit), Type 3 (4 bit), or Type 4 (8 bit) mode on the forward link, respectively; and Bit 3, Bit 4, or Bit 5 to select either a Type 2, Type 3, or Type 4 mode on the reverse link, respectively; with Bits 6 and 7 being reserved and generally set to zero at this time. The Bitmap Width and Height fields, here each being 2 bytes, specify the width and height of the bitmap, respectively, in pixels.

[00653] The Monochrome Capability field (1 byte) is used in one embodiment to specify the number of bits of resolution that can be displayed in a monochrome format. If a display cannot use a monochrome format then this value is set at zero. Bits 7 through 4 are reserved for future use and are, thus, set as zero. Bits 3 through 0 define the maximum number of bits of grayscale that can exist for each pixel. These four bits make it possible to specify values of 1 to 15 for each pixel. If the value is zero then monochrome format is not supported by the display.

[00654] The Bayer Capability field uses 1 byte to specify the number of bits of resolution, pixel group, and pixel order that can be transferred in Bayer format. If the client cannot use the Bayer format then this value is zero. The Bayer Capability field is composed of the following values: Bits 3 through 0 define the maximum number of bits of intensity that exist in each pixel, while Bits 5 through 4 define the pixel group pattern that is required, while Bits 8 through 6 define the pixel order that is required; with Bits 14 through 9 being reserved for future use and generally set to zero in the meantime. Bit 15, when set to a logic-one level indicates that the client can accept Bayer pixel data in either packed or unpacked format. If bit 15 is set to zero this indicates that the client can accept Bayer pixel data only in unpacked format.

[00655] The Color Map Capability field (3 bytes) in one embodiment specifies the maximum number of table items that exist in the color map table in the display. If the display cannot use the color map format then this value is set at zero.

[00656] The RGB Capability field (2 bytes) specifies the number of bits of resolution that can be displayed in RGB format. If the display cannot use the RGB format then this value is equal to zero. The RGB Capability word is composed of three separate unsigned values where: Bits 3 through 0 define the maximum number of bits of blue,

Bits 7 through 4 define the maximum number of bits of green, and Bits 11 through 8 define the maximum number of bits of red in each pixel. Currently, Bits 14 through 12 are reserved for future use and are generally set to zero. Bits 14 through 12 are reserved for future use and generally set to zero. Bit 15, when set to a logic-one level indicates that the client can accept RGB pixel data in either packed or unpacked format. If bit 15 is set to a logic-zero level, this indicates that the client can accept RGB pixel data only in unpacked format.

[00657] The Y Cr Cb Capability field (2 bytes) specifies the number of bits of resolution that can be displayed in Y Cr Cb format. If the display cannot use the Y Cr Cb format then this value is set equal to zero. The Y Cr Cb Capability word is composed of three separate unsigned values where: Bits 3 through 0 define the maximum number of bits in the Cb sample, Bits 7 through 4 define the maximum number of bits in the Cr sample, Bits 11 through 8 define the maximum number of bits in the Y sample, and Bits 15 through 12 are currently reserved for future use and are set to zero.

[00658] The Client Feature Capability field uses 4 bytes in one embodiment that contain a set of flags that indicate specific features in the client that are supported. A bit set to a logic-one level indicates the capability is supported, while a bit set to a logic-zero level indicates the capability is not supported. In one embodiment, the value for Bit 0 indicates whether or not Bitmap Block Transfer Packet (packet type 71) is supported. The value for Bits 1, 2, and 3 indicate whether or not Bitmap Area Fill Packet (packet type 72), Bitmap Pattern Fill Packet (packet type 73), or Read Frame Buffer Packet (packet type 74), respectively, are supported. The value for Bit 4 indicates whether or not the client has the capability to make one color transparent using the Transparent Color Enable Packet, while values for Bits 5 and 6 indicate if the client can accept audio data in unpacked or packed format, respectively, and the value for Bit 7 indicates whether or not the client can send a reverse-link video stream from a camera. The value for Bit 8 indicates whether or not the client has the ability to receive a full line of pixel data and ignore display addressing as specified by bit 5 of the Pixel Data Attributes field of the Video Stream Packet, and the client can also detect frame sync or end of video frame data using bit 15 of the Pixel Data Attributes Field.

[00659] The value of Bit 9 indicates whether or not the client has the ability to interpret the Request Specific Status Packet and respond with the Valid Status Reply List Packet. The client can indicate an ability to return additional status in the Valid Parameter Reply

List field of the Valid Status Reply List Packet as described above. The value of Bit 10 indicates whether or not the client has the ability to support display power state 01. The display power state is set using bits [3:2] of the Power State field of the Display Power State Packet described above. Display power state 01 is a state where the selected display is not illuminated and is consuming a minimum amount of power, if any, and the contents of the frame buffer are generally guaranteed to be retained during this state

[00660] The value for Bits 11 and 12 indicate when the client is communicating either with a pointing device and can send and receive Pointing Device Data Packets, or with a keyboard and can send and receive Keyboard Data Packets, respectively. The value for Bit 13 indicates whether or not the client has the ability to set one or more audio or video parameters by supporting the VCP Feature packets: Request VCP Feature Packet, VCP Feature Reply Packet, Set VCP Feature Packet, Request Valid Parameter Packet, and Valid Parameter Reply Packet. The value for Bit 14 indicates whether or not the client has the ability to write pixel data into the offline display frame buffer, which is illustrated in FIG. 88A. If this bit is set to a logic-one level then the Display Update Bits (bits 7 and 6 of the Pixel Data Attributes field of the Video Stream Packet) may be set to the values '01'.

[00661] The value for Bit 15 indicates when the client has the ability to write pixel data into only the display frame buffer currently being used to refresh the display image, which is illustrated in FIG. 88B. If this bit is set to a logic-one then the Display Update Bits (bits 7 and 6 of the Pixel Data Attributes field of the Video Stream Packet) may be set to the values '00'. The value for Bit 16 indicates when the client has the ability to write pixel data from a single Video Stream Packet into all display frame buffers, which is illustrated in FIG. 88C. If this bit is set equal to a logic-one level then the Display Update Bits (bits 7 and 6 of the Pixel Data Attributes field of the Video Stream Packet) may be set to the value '11'.

[00662] In one embodiment, the value for Bit 17 indicates when a client has the ability to respond to the Request Specific Status Packet, the value for Bit 18 indicates when the client has the ability to respond to the Round Trip Delay Measurement Packet, and the value for Bit 19 indicates when the client has the ability to the Forward Link Skew Calibration Packet. In one embodiment, the value for Bit 20 indicates when the client has the ability to respond to the Display Power State Packet.

- [00663] In one embodiment, the value for Bit 21 indicates when the client has the ability to use the Raster Operation field of the Block Transfer Packet (packet type 71), the Bitmap Area Fill Packet (packet type 72), and the Bitmap Pattern Fill Packet (packet type 73) if those packets are supported by the client as specified by bits 0, 1, and 2 of this field. In one embodiment, if bit 21 has a logic-zero level or value, and the packets are supported, then the client does not have the ability to use the Raster Operation field and the client only has the ability to copy or write to pixel locations specified by these packets.
- [00664] The value for Bit 22 indicates whether or not the client has the ability to respond to the Register Access Packet. Bits 23 through 31 are currently reserved for future use or alternative designations useful for system designers, and are generally set equal to a zero value or a logic-zero level.
- [00665] The Display Video Frame Rate Capability field (1 byte) specifies the maximum video frame update capability of the display in frames per second. A host may choose to update the image at a slower rate than the value specified in this field.
- [00666] The Audio Buffer Depth field (2 bytes) specifies the depth of the elastic buffer in a Display which is dedicated to each audio stream.
- [00667] The Audio Channel Capability field (2 bytes) contains a group of flags that indicate which audio channels are supported by the client or client connected device. A bit set to one indicates the channel is supported, and a bit set to zero indicates that channel is not supported. The Bit positions are assigned to the different channels, for example Bit positions 0, 1, 2, 3, 4, 5, 6, and 7 in one embodiment, indicate the left front, right front, left rear, right rear, front center, sub-woofer, surround left, and surround right channels, respectively. Bits 8 through 14 are currently reserved for future use, and are generally set to zero. In one embodiment Bit 15 is used to indicate if the client provides support for the Forward Audio Channel Enable Packet. If this is the case, Bit 15 set to a logic-one level. If, however, the client is not capable of disabling audio channels as a result of the Forward Audio Channel Enable Packet or if the client does not support any audio capability, then this bit is set to a logic-zero level or value.
- [00668] A 2-byte Audio Sample Rate Capability field, for the forward link, contains a set of flags to indicate the audio sample rate capability of the client device. Bit positions are assigned to the different rates accordingly, such as Bits 0, 1, 2, 3, 4, 5, 6, 7, and 8 being assigned to 8,000, 16,000, 24,000, 32,000, 40,000, 48,000, 11,025, 22,050, and

44,100 samples per second (SPS), respectively, with Bits 9 through 15 being reserved for future or alternative rate uses, as desired, so they are currently set to '0'. Setting a bit value for one of these bits to '1' indicates that that particular sample rate is supported, and setting the bit to '0' indicates that that sample rate is not supported.

[00669] The Minimum Sub-frame Rate field (2 bytes) specifies the minimum sub-frame rate in frames per second. The minimum sub-frame rate keeps the client status update rate sufficient to read certain sensors or pointing devices in the client.

[00670] A 2-byte Mic Sample Rate Capability field, for the reverse link, contains a set of flags that indicate the audio sample rate capability of a microphone in the client device. For purposes of the MDDI, a client device microphone is configured to minimally support at least an 8,000 sample per second rate. Bit positions for this field are assigned to the different rates with bit positions 0, 1, 2, 3, 4, 5, 6, 7, and 8, for example, being used to represent 8,000, 16,000, 24,000, 32,000, 40,000, 48,000, 11,025, 22,050, and 44,100 samples per second (SPS), respectively, with Bits 9 through 15 being reserved for future or alternative rate uses, as desired, so they are currently set to '0'. Setting a bit value for one of these bits to '1' indicates that that particular sample rate is supported, and setting the bit to '0' indicates that that sample rate is not supported. If no microphone is connected then each of the Mic Sample Rate Capability bits are set equal to zero.

[00671] The Keyboard Data Format field (here 1 byte) specifies whether or not a keyboard is connected to the client system and the type of keyboard that is connected. In one embodiment, the value established by Bits 6 through 0 is used to define the type of keyboard that is connected. If the value is zero (0) then the keyboard type is considered as unknown. For a value of 1, the keyboard data format is considered to be a standard PS-2 style. Currently values in the range of 2 through 125 are not in use, being reserved for use of system designers and interface incorporators or product developers to define specific keyboard or input devices for use with the MDDI and corresponding clients or hosts. A value of 126 is used to indicate that the keyboard data format is user-defined, while a value of 127 is used to indicate that a keyboard cannot be connected to this client. In addition, Bit 7 can be used to indicate whether or not the keyboard can communicate with the client. The intended use of this bit is to indicate when the keyboard can communicate with the client using a wireless link. Bit 7 would be set to a zero level if bits 6 through 0 indicate that a keyboard cannot be connected to the client.

Therefore, for one embodiment, when the value of Bit 7 is 0, the keyboard and client cannot communicate, while if the value of Bit 7 is 1, the keyboard and client have acknowledged that they can communicate with each other.

[00672] The Pointing Device Data Format field (here 1 byte) specifies whether or not a pointing device is connected to the client system and the type of pointing device that is connected. In one embodiment, the value established by Bits 6 through 0 is used to define the type of pointing device that is connected. If the value is zero (0) then the pointing device type is considered as unknown. For a value of 1, the pointing device data format is considered to be a standard PS-2 style. Currently values in the range of 2 through 125 are not in use, being reserved for use of system designers and interface incorporators or product developers to define specific pointing device or input devices for use with the MDDI and corresponding clients or hosts. A value of 126 is used to indicate that the pointing device data format is user-defined, while a value of 127 is used to indicate that a pointing device cannot be connected to this client. In addition, Bit 7 can be used to indicate whether or not the pointing device can communicate with the client. The intended use of this bit is to indicate when the keyboard can communicate with the client using a wireless link. Bit 7 would be set to a zero level if bits 6 through 0 indicate that a pointing device cannot be connected to the client. Therefore, for one embodiment, when the value of Bit 7 is 0, the pointing device and client cannot communicate, while if the value of Bit 7 is 1, the pointing device and client have acknowledged that they can communicate with each other.

[00673] The Content Protection Type field (2 bytes) contains a set of flags that indicate the type of digital content protection that is supported by the Display. Currently, bit position 0 is used to indicate when DTCP is supported and bit position 1 is used to indicate when HDCP is supported, with bit positions 2 through 15 being reserved for use with other protection schemes as desired or available, so they are currently set to zero.

[00674] The Mfr Name field (here 2 bytes) contains the EISA 3-character ID of the manufacturer, packed into three 5-bit characters in the same manner as in the VESA EDID specification. The character 'A' is represented as 00001 binary, the character 'Z' is represented as 11010 binary, and all letters between 'A' and 'Z' are represented as sequential binary values that correspond to the alphabetic sequence between 'A' and 'Z'. The most significant bit of the Mfr Name field is unused and is generally set to

logic-zero for now until a use is made in the future implementations. For example, a manufacturer represented by the string "XYZ" would have a Mfr Name value of 0x633a. If this field is not supported by the client it is generally set to zero. Product Code field uses 2 bytes to contain a product code assigned by the display manufacturer. If this field is not supported by the client it is generally set to zero.

[00675] Reserved 1, Reserved 2, and Reserved 3 fields (here 2 bytes each) are reserved for future use in imparting information. All bits in these field are generally be set to a logic-zero level. The purpose of such fields is currently to cause all subsequent 2 byte fields to align to a 16-bit word address and cause 4-byte fields to align to a 32-bit word address.

[00676] The Serial Number field uses 4 bytes in this embodiment to specify the serial number of the display in numeric form. If this field is not supported by the client it is generally set to zero. The Week of Manufacture field uses 1 byte to define the week of manufacture of the display. This value is typically in the range of 1 to 53 if it is supported by the client. If this field is not supported by the client it is set to zero. The Year of Manufacture field is 1 byte that defines the year of manufacture of the display. This value is an offset from the year 1990. Years in the range of 1991 to 2245 can be expressed by this field. Example: the year 2003 corresponds to a Year of Manufacture value of 13. If this field is not supported by the client it is set to zero.

[00677] The CRC field (here 2 bytes) contains a 16-bit CRC of all bytes in the packet including the Packet Length.

G. For Client Request and Status Packets

[00678] The Reverse Link Request field (3 byte) specifies the number of bytes the client needs in the reverse link in the next sub-frame to send information to the host.

[00679] The CRC Error Count field (1 byte) indicates how many CRC errors have occurred since the beginning of the media-frame. The CRC count is reset when a sub-frame header packet with a Sub-frame Count of zero is sent. If the actual number of CRC errors exceeds 255 then this value generally saturates at 255.

[00680] The Capability Change field uses 1 byte to indicate a change in the capability of the client. This could occur if a user connects a peripheral device such as a microphone, keyboard, or display, or for some other reason. When Bits[7:0] are equal to 0, then the capability has not changed since the last Client Capability Packet was sent. However,

when Bits[7:0] are equal to 1 to 255, the capability has changed. The Client Capability Packet is examined to determine the new display characteristics.

[00681] The Client Busy Flags field uses 2 bytes to indicate that the client is performing a specific function and is not ready to yet accept another packet related to that function. A bit set to a logic-one level or value indicates that the particular function is currently being performed by the client and that the related function in the client is busy. If the related function in the client is ready, the bit is set to a logic-zero. The client should return a busy status (bit set to one) for all functions that are not supported in the client.

[00682] In one embodiment, these bytes are interpreted according to the following relationships. If Bit 0 is a '1' then the bitmap block transfer function is busy, while if Bit 1 is a '1', then a bitmap area fill function is busy, and if Bit 2 is a '1', then a bitmap pattern fill function is busy. At the same time, if Bit 3 is a '1' then the graphics subsystem is busy performing an operation that requires use of the frame buffer in the client. Other graphics functions that require use of the frame buffer may not begin until this bit is set to logic-one.

[00683] Currently, Bits 4 through 15 remain reserved for future use and are generally set to a logic-one level or state to indicate a busy status in a case these bits are assigned in the future.

H. For Bit Block Transfer Packets

[00684] The Window Upper Left Coordinate X Value and Y Value fields use 2 bytes each to specify the X and Y value of the coordinates of the upper left corner of the window to be moved. The Window Width and Height fields use 2 bytes each to specify the width and height of the window to be moved. The Window X Movement and Y Movement fields use 2 bytes each to specify the number of pixels that the window is to be moved horizontally and vertically, respectively. Typically, these coordinates are configured such that positive values for X cause the window to be moved to the right, and negative values cause movement to the left, while positive values for Y cause the window to be moved down, and negative values cause upward movement.

I. For Bitmap Area Fill Packets

[00685] Window Upper Left Coordinate X Value and Y Value fields use 2 bytes each to specify the X and Y value of the coordinates of the upper left corner of the window to

be filled. The Window Width and Height fields (2 bytes each) specify the width and height of the window to be filled. The Video Data Format Descriptor field (2 bytes) specifies the format of the Pixel Area Fill Value. The format is the same as the same field in the Video Stream Packet. The Pixel Area Fill Value field (4 bytes) contains the pixel value to be filled into the window specified by the fields discussed above. The format of this pixel is specified in the Video Data Format Descriptor field.

J. For Bitmap Pattern Fill Packets

[00686] Window Upper Left Coordinate X Value and Y Value fields use 2 bytes each to specify the X and Y value of the coordinates of the upper left corner of the window to be filled. The Window Width and Height fields (2 bytes each) specify the width and height of the window to be filled. The Pattern Width and Pattern Height fields (2 bytes each) specify the width and height, respectively, of the fill pattern. The Horizontal Pattern Offset field (2 bytes) specifies a horizontal offset of the pixel data pattern from the left edge of the specified window to be filled. The value being specified is to be less than the value in the Pattern Width Field. The Vertical Pattern Offset field (2 bytes) specifies a vertical offset of the pixel data pattern from the top edge of the specified window to be filled. The value being specified is to be less than the value in the Pattern Height field.

[00687] The 2-byte Video Data Format Descriptor field specifies the format of the Pixel Area Fill Value. FIG. 11 illustrates how the Video Data Format Descriptor is coded. The format is the same as the same field in the Video Stream Packet.

[00688] The Parameter CRC field (2 bytes) contains a CRC of all bytes from the Packet Length to the Video Format Descriptor. If this CRC fails to check then the entire packet is discarded. The Pattern Pixel Data field contains raw video information that specifies the fill pattern in the format specified by the Video Data Format Descriptor. Data is packed into bytes, and the first pixel of each row is to be byte-aligned. The fill pattern data is transmitted a row at a time. The Pattern Pixel Data CRC field (2 bytes) contains a CRC of only the Pattern Pixel Data. If this CRC fails to check then the Pattern Pixel Data can still be used but the CRC error count is incremented.

K. Communication Link Data Channel Packets

- [00689] The Parameter CRC field (2 bytes) contain a 16-bit CRC of all bytes from the Packet Length to the Packet Type. If this CRC fails to check then the entire packet is discarded.
- [00690] The Communication Link Data field contains the raw data from the communication channel. This data is simply passed on to the computing device in the display.
- [00691] The Communication Link Data CRC field (2 bytes) contains a 16-bit CRC of only the Communication Link Data. If this CRC fails to check then the Communication Link Data is still used or useful, but the CRC error count is incremented.

L. For Forward Audio Channel Enable Packets

- [00692] The Audio Channel Enable Mask field (1 byte) contains a group of flags that indicate which audio channels are to be enabled in a client. A bit set to one enables the corresponding channel, and a bit set to zero disables the corresponding channel. Bits 0 through 5 designate channels 0 through 5 which address left front, right front, left rear, right rear, front center, and sub-woofer channels, respectively. Bits 6 and 7 are reserved for future use, and in the mean time are generally set equal to zero.

M. For Reverse Audio Sample Rate Packets

- [00693] The Audio Sample Rate field (1 byte) specifies the digital audio sample rate. The values for this field are assigned to the different rates with values of 0, 1, 2, 3, 4, 5, 6, 7, and 8 being used to designate 8,000, 16,000, 24,000, 32,000, 40,000, 48,000, 11,025, 22,050, and 44,100 samples per second (SPS), respectively, with values of 9 through 254 being reserved for use with alternative rates, as desired, so they are currently set to '0'. A value of 255 is used to disable the reverse-link audio stream.
- [00694] The Sample Format field (1 byte) specifies the format of the digital audio samples. When Bits[1:0] are equal to '0', the digital audio samples are in linear format, when they are equal to 1, the digital audio samples are in μ -Law format, and when they are equal to 2, the digital audio samples are in A-Law format. Bits[7:2] are reserved for alternate use in designating audio formats, as desired, and are generally set equal to zero.

N. For The Digital Content Protection Overhead Packets

- [00695] The Content Protection Type field (1 byte) specifies the digital content protection method that is used. A value of '0' indicates Digital Transmission Content Protection (DTCP) while a value of 1 indicates High-bandwidth Digital Content Protection System (HDCP). The value range of 2 through 255 is not currently specified but is reserved for use with alternative protection schemes as desired. The Content Protection Overhead Messages field is a variable length field containing content protection messages sent between the host and client.

O. For The Transparent Color Enable Packets

- [00696] The Transparent Color Enable field (1 byte) specifies when transparent color mode is enabled or disabled. If Bit 0 is equal to 0 then transparent color mode is disabled, if it is equal to 1 then transparent color mode is enabled and the transparent color is specified by the following two parameters. Bits 1 through 7 of this byte are reserved for future use and are typically set equal to zero.
- [00697] The Video Data Format Descriptor field (2 bytes) specifies the format of the Pixel Area Fill Value. FIG. 11 illustrates how the Video Data Format Descriptor is coded. The format is generally the same as the same field in the Video Stream Packet.
- [00698] The Pixel Area Fill Value field uses 4 bytes allocated for the pixel value to be filled into the window specified above. The format of this pixel is specified in the Video Data Format Descriptor field.

P. For The Round Trip Delay Measurement Packets

- [00699] The 2-byte Packet Length field specifies the total number of bytes in the packet not including the packet length field, and in one embodiment is selected to have a fixed length of 159. The 2-byte Packet Type field that identifies this packet type with a value of 82, identifying a packet as a Round Trip Delay Measurement Packet. The Client ID field, as before is reserved for future use as a Client ID, and is generally set to zero.
- [00700] In one embodiment, the Parameter CRC field (2 bytes) contains a 16-bit CRC of all bytes from the Packet Length to the Packet Type. If this CRC fails to check then the entire packet is discarded.
- [00701] The Guard Time 1 field (here 64 bytes) is used to allow the MDDI Data line drivers in the client to enable before the line drivers in the host are disabled. The client

enables its MDDI_Data line drivers during bit 0 of Guard Time 1 and the host disables its line drivers so as to be completely disabled prior to the last bit of Guard Time 1. The host and client both drive a logic-zero level during Guard Time 1 when they are not disabled. Another purpose of this field is to ensure that all MDDI_Data signals are at a logic-zero level for a sufficient time to allow the client to begin recovering a clock or clock signal using only MDDI_Stb prior to disabling the host's line drivers.

[00702] The Measurement Period field is a 64 byte window used to allow the client to respond with two bytes of 0xff, and 30 bytes of 0x00 at half the data rate used on the forward link. This data rate corresponds to a Reverse Link Rate Divisor of 1. The client returns this response immediately at the time it perceives as being the beginning of the Measurement Period. This response from the client will be received at a host at precisely the round trip delay of the link plus logic delay in the client after the beginning of the first bit of the Measurement Period at the host.

[00703] The All Zero 1 field (2 bytes) contains zeroes to allow the MDDI_Data line drivers in the host and client to overlap so that MDDI_Data is always driven. The host enables MDDI_Data line drivers during bit 0 of the All Zero 1 field, and the client also continues to drive the signal to a logic-zero level as it did at the end of the Measurement Period.

[00704] The value in the Guard Time 2 field (64 bytes) allows overlap of the Measurement Period driven by the client when the round trip delay is at the maximum amount that can be measured in the Measurement Period. The Client disables its line drivers during bit 0 of Guard Time 2 and the Host enables its line drivers immediately after the last bit of Guard Time 2. The host and client both drive a logic-zero level during Guard Time 2 when they are not disabled. Another purpose of this field is to ensure that all MDDI_Data signals are at a logic-zero level for a sufficient time to allow the client to begin recovering a clock signal using both MDDI_Data0 and MDDI_Stb after enabling the line drivers for a host.

Q. For The Forward Link Skew Calibration Packets

[00705] In one embodiment, the Parameter CRC field (2 bytes) contains a 16-bit CRC of all bytes from the Packet Length to the Packet Type. If this CRC fails to check then the entire packet is discarded.

[00706] The All Zero 1 field uses 8 bytes to ensure that there will be an transitions on the MDDI_Stb at the beginning of the Calibration Data Sequence field. Generally, these bytes employ 8-bit unsigned integers equal to zero. It also provides for sufficient time for the client core logic to change the mode of the clock recovery circuit from using the XOR of MDDI_0 and MDDI_Stb, to simply using MDDI_Stb or the MDDI_Stb signal as the recovered clock.

[00707] The Calibration Data Sequence field contains a data sequence that causes the MDDI_Data signals to toggle at every data period. The length of the Calibration Data Sequence field is determined by the interface being used on the forward link. During the processing of the Calibration Data Sequence, the MDDI host controller sets all MDDI_Data signals equal to the strobe signal. The client clock recovery circuit should use only MDDI_Stb rather than MDDI_Stb XOR MDDI_Data0 to recover the data clock while the Calibration Data Sequence field is being received by the client. Depending on the exact phase of the MDDI_Stb signal at the beginning of the Calibration Data Sequence field, the Calibration Data Sequence will generally be one of the following based on the interface Type being used when this packet is sent:

Type 1 – (64 byte data sequence) 0xaa, 0xaa ... or 0x55, 0x55...

Type 2 – (128 byte data sequence) 0xcc, 0xcc ... or 0x33, 0x33...

Type 3 – (256 byte data sequence) 0xf0, 0xf0 ... or 0x0f, 0x0f ...

Type 4 – (512 byte data sequence) 0xff, 0x00, 0xff, 0x00 ... or 0x00, 0xff, 0x00, 0xff ...

[00708] The All Zero 2 field uses 8 bytes to provide sufficient time for the client core logic to change the mode of the clock recovery circuit back to an original state, from using the MDDI_Stb signal as the recovered clock to using the XOR of MDDI_0 and MDDI_Stb. Generally, these bytes employ 8-bit unsigned integers equal to zero.

[00709] An example of the possible MDDI_Data and MDDI_Stb waveforms for both the Type 1 and Type 2 Interfaces are shown in FIGs. 62A and 62B, respectively.

XIX. Conclusion

[00710] While various embodiments of the present invention have been described above, it should be understood that they have been presented by way of example only, and not limitation. Thus, the breadth and scope of the present invention should not be limited

by any of the above-described exemplary embodiments, but should be defined only in accordance with the following claims and their equivalents.

[00711] **WHAT IS CLAIMED IS:**

CLAIMS

1. A digital data interface for transferring digital presentation data at a high rate between a host device and a client device over a communication path comprising:
a plurality of packet structures linked together to form a communication protocol for communicating a pre-selected set of digital control and presentation data between a host and a client over said communication path; and
at least one link controller residing in said host device coupled to said client through said communications path, being configured to generate, transmit, and receive packets forming said communications protocol, and to form presentation data into one or more types of data packets.
2. The interface of Claim 1 further comprising said packets grouped together within media frames that are communicated between said host and client having a pre-defined fixed length with a pre-determined number of said packets having differing and variable lengths.
3. The interface of Claim 1 further comprising a Sub-frame Header Packet positioned at the beginning of transfers of packets from said host.
4. The interface of Claim 1 further comprising a Display Power State Packet to place specific client related hardware into a low power state.
5. The interface of Claim 1 wherein said link controller is a host link controller and further comprising at least one client link controller residing in said client device coupled to said host through said communications path, being configured to generate, transmit, and receive packets forming said communications protocol, and to form digital presentation data into one or more types of data packets.
6. The interface of Claim 2 further comprising:
a plurality of transfer modes, each allowing the transfer of different maximum numbers of bits of data in parallel over a given time period, with each mode selectable by negotiation between said host and client link drivers; and

wherein said transfer modes are dynamically adjustable between said modes during transfer of data.

7. The interface of Claim 1 further comprising a Link Shutdown type packet for transmission by said host to said client to terminate the transfer of data in either direction over said communication path.

8. The interface of Claim 1 further comprising means for said client to wake up said host from a hibernation state.

9. A method of transferring digital data at a high rate between a host device and a client device over a communication path for presentation to a user, comprising:

generating one or more of a plurality of pre-defined packet structures and linking them together to form a pre-defined communication protocol;

communicating a pre-selected set of digital control and presentation data between said host and said client devices over said communication path using said communication protocol;

coupling at least one host link controller residing in said host device to said client device through said communications path, the host link controller being configured to generate, transmit, and receive packets forming said communications protocol, and to form digital presentation data into one or more types of data packets; and

transferring data in the form of packets over said communications path using said link controllers.

10. The method of Claim 9 further comprising grouping said packets together within media frames for communication between said host and client, the media frames having a pre-defined fixed length with a pre-determined number of said packets having differing and variable lengths.

11. The method of Claim 9 further comprising commencing transfer of packets from said host with a Sub-frame Header type packet.

12. The method of Claim 9 further comprising generating, transmitting, and receiving packets forming said communications protocol through at least one client link controller residing in said client device coupled to said host device through said communications path.

13. The method of Claim 12 further comprising:

negotiating between host and client link drivers the use of one of a plurality of transfer modes in each direction, each allowing the transfer of different maximum numbers of bits of data in parallel over a given time period; and

dynamically adjusting between said transfer modes during transfer of data.

14. The method of Claim 9 further comprising terminating the transfer of data in either direction over said communication path using a Link Shutdown type packet for transmission by said host to said client.

15. The method of Claim 9 further comprising waking up said host from a hibernation state by communication with said client.

16. The method of Claim 9 further comprising generating a Display Power State Packet to place specific client hardware into a low power state.

17. Apparatus for transferring digital data at a high rate between a host device and a client device over a communication path for presentation to a user, comprising:

at least one host link controller disposed in said host device for generating one or more of a plurality of pre-defined packet structures and linking them together to form a pre-defined communication protocol, and for communicating a pre-selected set of digital control and presentation data between said host and said client devices over said communication path using said communication protocol;

at least one client controller disposed in said client device and coupled to said host link controller through said communications path; and

each link controller being configured to generate, transmit, and receive packets forming said communications protocol, and to form digital presentation data into one or more types of data packets.

18. The apparatus of Claim 17 wherein said host controller comprises a state machine.
19. The apparatus of Claim 17 wherein said host controller comprises a general purpose signal processor.
20. The apparatus of Claim 17 further comprising a Sub-frame Header type packet at the commencing of transfer of packets from said host.
21. The apparatus of Claim 17 wherein said host controller comprises one or more differential line drivers; and said client receiver comprises one or more differential line receivers coupled to said communication path.
22. The apparatus of Claim 17 wherein said host and client link controllers are configured to use of one of a plurality of transfer modes in each direction, each allowing the transfer of different maximum numbers of bits of data in parallel over a given time period; and being capable of being dynamically adjusting between said transfer modes during transfer of data.
23. The apparatus of Claim 17 wherein said host controller is configured to transmit a Link Shutdown type packet to said client means for terminating the transfer of data in either direction over said communication path.
24. The apparatus of Claim 17 wherein said host controller is configured to generate a Display Power State Packet to place specific display controller hardware into a low power state.

25. For use in an electronic system for transferring digital data at a high rate between a host device and a client device over a communication path for presentation to a user, a computer program product comprising:

a computer usable medium having computer readable program code means embodied in said medium for causing an application program to execute on the computer system, said computer readable program code means comprising:

a computer readable first program code means for causing the computer system to generate one or more of a plurality of pre-defined packet structures and link them together to form a pre-defined communication protocol;

a computer readable second program code means for causing the computer system to communicate a pre-selected set of digital control and presentation data between said host and said client devices over said communication path using said communication protocol;

a computer readable third program code means for causing the computer system to couple at least one host link controller disposed in said host device to at least one client controller disposed in said client device through said communications path, the link controllers being configured to generate, transmit, and receive packets forming said communications protocol, and to form digital presentation data into one or more types of data packets; and

a computer readable fourth program code means for causing the computer system to transfer data in the form of packets over said communications path using said link controllers.

26. The computer program product of Claim 25, comprising:

a computer readable program code means for causing the computer system to generate a pre-defined packet structures and link them together to form a pre-defined communication protocol;

27. Apparatus for transferring digital data at a high rate between a host device and a client device over a communication path for presentation to a user, comprising:

means for generating one or more of a plurality of pre-defined packet structures and linking them together to form a pre-defined communication protocol;

means for communicating a pre-selected set of digital control and presentation data between said host and said client devices over said communication path using said communication protocol;

means for coupling at least two link controllers together through said communications path, one in each of said host and client and each being configured to generate, transmit, and receive packets forming said communications protocol, and to form digital presentation data into one or more types of data packets; and

means for transferring data in the form of packets over said communications path using said link controllers.

28. The apparatus of Claim 27 further comprising means for commencing transfer of packets from said host with a Sub-frame Header type packet.

29. The apparatus of Claim 27 further comprising means for comprising generating a Display Power State Packet configured to place specific hardware into a low power state.

30. The apparatus of Claim 27 further comprising means for requesting display capabilities information from the client by a host link controller so as to determine what type of data and data rates said client is capable of accommodating through said interface.

31. A processor for use in an electronic system for transferring digital data at a high rate between a host device and a client device over a communication path, the processor configured to generate one or more of a plurality of pre-defined packet structures and link them together to form a pre-defined communication protocol; to form digital presentation data into one or more types of data packets; communicate a pre-selected set of digital control and presentation data between said host and said client devices over said communication path using said communication protocol; and transfer data in the form of packets over said communications path.

32. A state machine for use in obtaining synchronization in an electronic system transferring digital data at a high rate between a host device and a client device over a

communication path, the state machine configured to have at least one Async Frames State synchronization state, at least two Acquiring Sync States synchronization states, and at least three In-Sync States synchronization states.

33. A state machine for use in obtaining synchronization in an electronic system transferring digital data at a high rate between a host device and a client device over a communication path, the state machine configured to have at least one Acquiring Sync States synchronization states, and at least two In-Sync States synchronization states.

1/86




FIG. 1A

2/86




FIG. 1B

3/86




FIG. 2B




FIG. 2A

4/86

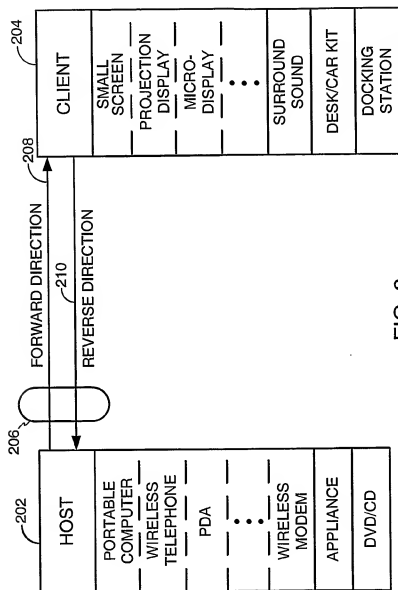


FIG. 3



FIG. 4

5/86




FIG. 5

6/86




FIG. 6

7/86

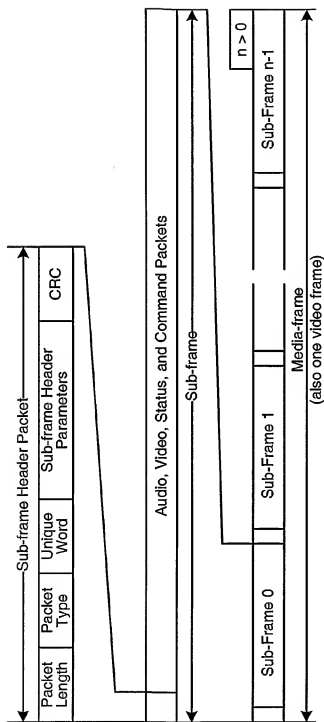




FIG. 7



FIG. 8



Packet Length	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes
Packet Type = 16	bClient ID	Video Data Format Descriptor	Pixel Data Attributes	X Left Edge	Y Top Edge	X Right Edge	Y Bottom Edge				

FIG. 11

X Start	Y Start	Pixel Count	Parameter CRC	Pixel Data	Pixel Data CRC
2 bytes	2 bytes	2 bytes	2 bytes	Packet Length - 26 bytes	2 bytes

Video Data Format Descriptor - Monochrome

FIG. 12A	0	0	0	P	not used = 0	number of bits per pixel
	15	14	13	12	11	4 3 0

Video Data Format Descriptor - Color Using Color Map

FIG. 12B	0	0	1	P	not used = 0	number of bits per pixel
	15	14	13	12	11	6 5 0

Video Data Format Descriptor - Color, Raw RGB

FIG. 12C	0	1	0	P	number of bits of red per pixel	number of bits of green per pixel	number of bits of blue per pixel
	15	14	13	12	11	8 7	4 3 0

Video Data Format Descriptor - Color Component Video, YCbCr

FIG. 12D	0	1	1	P	number of bits of Y per pixel	number of bits of Cb per pixel	number of bits of Cr per pixel
	15	14	13	12	11	8 7	4 3 0

Video Data Format Descriptor - Bayer

FIG. 12E	1	0	0	P	not used = 0	Pixel order	Pixel Pattern	number of bits per pixel
	15	14	13	12	11	9 8	6 5 4 3	0

10/86






FIG. 13

11/86

Audio Stream Packet

Packet Length	Packet Type = 32	bClient ID	Audio Channel ID	Reserved 1	Audio Sample Count
2 bytes	2 bytes	2 bytes	1 byte	1 byte	2 bytes
Bits per Sample and Packing					
1 byte		Audio Sample Rate	Parameter CRC	Digital Audio Data	Audio Data CRC
1 byte		1 byte	2 bytes	Packet Length - 14 bytes	2 bytes

FIG. 14

Packet Length	Packet Type = 56 to 63	bClient ID	Stream Parameters	Parameter CRC	Stream Data	Stream Data CRC
2 bytes	2 bytes	2 bytes	2 bytes		2 bytes	

FIG. 16

Packet Length	hClient ID	Color Map Item Count	Color Map Offset	Color Map Item Size	Reserved 1	Parameter CRC	Color Map Data	Data CRC
2 bytes	2 bytes	2 bytes	4 bytes	1 byte	1 byte	2 bytes	(Packet Length - 16) bytes	2 bytes

FIG. 17

12/86

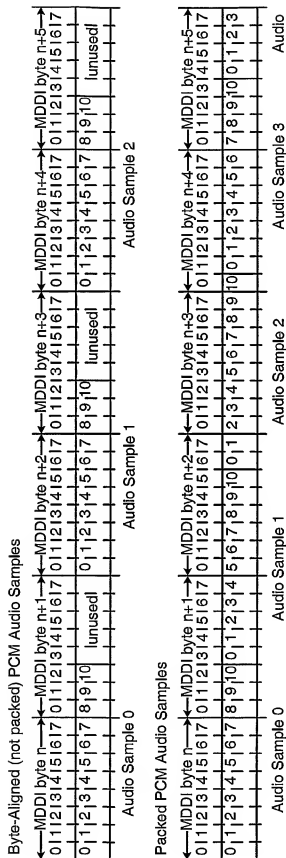


FIG. 15

13/86

Reverse Link Encapsulation Packet

Packet Length	Packet Type = 65	Reverse Link flags	Reverse Rate Divisor	Turn-Around 1 Length	Turn-Around 2 Length
2 bytes	2 bytes	1 byte	1 byte	1 byte	1 byte

Parameter CRC	All Zero 1	Turn-Around 1	reverse data packets	Turn-Around 2	All Zero 2
2 bytes	8 bytes	x bytes	(Packet Length - x - y - 26) bytes	y bytes	8 bytes

FIG. 18

Client Capability Packet

Packet Length	Packet Type = 66	cClient ID	Protocol Version	Min Protocol Version	Data Rate Capability	Interface Type Capability
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	1 byte

Number of Alt Displays	Reserved 1	Bitmap Width	Bitmap Height	Display Window Width	Display Window Height	Color Map Size
1 byte	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	4 bytes

Color Map RGB Width	RGB Capability	Monochrome Capability	Reserved 2	Y Cb Cr Capability	Bayer Capability	Reserved 3
2 bytes	2 bytes	1 byte	1 byte	2 bytes	2 bytes	2 bytes

Client Feature Capability	Max Video Frame Rate	Min Video Frame Rate	Min Sub-Frame rate	Audio Buffer Depth	Audio Channel Rate Capability	Audio Sample Rate Capability
4 bytes	1 byte	1 byte	2 bytes	2 bytes	2 bytes	2 bytes

Audio Sample Resolution	Mic Sample Resolution	Mic Sample Rate Capability	Keyboard Data Format	Poining Device Data Format	Content Protection Type	Mfr Name
1 byte	1 byte	2 bytes	1 byte	1 byte	2 bytes	2 bytes

Product Code	Reserved 4	Serial Number	Week of Mfr	Year of Mfr	CRC
2 bytes	2 bytes	4 bytes	1 byte	1 byte	2 bytes

FIG. 19

14/86

Packet Length	Packet Type = 67	bClient ID	Keyboard Data Format	Keyboard Data	CRC
2 bytes	2 bytes	2 bytes	2 bytes	(Packet_Length - 8) bytes	2 bytes

FIG. 20

Packet Length	Packet Type = 68	bClient ID	Pointing Device Format	Pointing Device Data	CRC
2 bytes	2 bytes	2 bytes	2 bytes	(Packet_Length - 8) bytes	2 bytes

FIG. 21

15/86

Packet Length	Packet Type = 69	CRC	All Zeros
2 bytes	2 byte	2 bytes	16 bytes

FIG. 22

Packet Length	Packet Type = 70	cClient ID	Reverse Link Request	CRC Error Count	Capability Change	Graphics Busy flags	CRC
2 bytes	2 bytes	2 bytes	2 bytes	1 byte	1 byte	2 bytes	2 bytes

FIG. 23

Bitmap Block Transfer Packet

Packet Length	Packet Type = 71	hClient ID	Pixel Data Attributes	Raster Operation	Upper Left X Value
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

Upper Left Y Value	Window Width	Window Height	Window X Movement	Window Y Movement	CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

FIG. 24

16/86

Bitmap Area Fill Packet

Packet Length	Packet Type = 72	hClient ID	Video Data Format Descriptor	Pixel Data Attributes	Pixel Area Fill Value
---------------	------------------	------------	------------------------------	-----------------------	-----------------------

2 bytes 2 bytes 2 bytes 2 bytes 2 bytes 4 bytes

Upper Left X Value	Upper Left Y Value	Window Width	Window Height	CRC
--------------------	--------------------	--------------	---------------	-----

2 bytes 2 bytes 2 bytes 2 bytes 2 bytes

FIG. 25**Bitmap Pattern Fill Packet**

Packet Length	Packet Type = 73	hClient ID	Video Data Format Descriptor	Pixel Data Attributes	Upper Left X Value	Upper Left Y Value	Window Width	Window Height
---------------	------------------	------------	------------------------------	-----------------------	--------------------	--------------------	--------------	---------------

2 bytes 2 bytes 2 bytes 2 bytes 2 bytes 2 bytes 2 bytes 2 bytes 2 bytes

Pattern Width	Pattern Height	Horizontal Pattern Offset	Vertical Pattern Offset	Parameter CRC	Pattern Pixel Data	Pixel Data CRC
---------------	----------------	---------------------------	-------------------------	---------------	--------------------	----------------

2 bytes 2 bytes 2 bytes 2 bytes 2 bytes length - 22 bytes 2 bytes

FIG. 26

17/86




FIG. 27

Display Power State Packet




FIG. 28

Perform Type Handoff Packet

Packet Length	Packet Type = 77	Interface Type	Reserved 1	Delay Filler	CRC
2 bytes	2 bytes	1 byte	1 byte	Packet Length - 6 bytes	2 bytes

FIG. 29

18/86

Packet Length	Packet Type = 78	hClient ID	Audio Channel Enable Mask	CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

FIG. 30

Packet Length	Packet Type = 79	hClient ID	Audio Sample Rate	Reserved 1	CRC
2 bytes	2 bytes	2 bytes	1 byte	1 byte	2 bytes

FIG. 31

19/86

Packet Length 2 bytes	Packet Type = 80 2 bytes	bClient ID 2 bytes	Content Protection Type 2 bytes	Content Protection Overhead Messages Packet Length - 8 bytes	CRC 2 bytes
--------------------------	-----------------------------	-----------------------	------------------------------------	---	----------------

FIG. 32

Packet Length 2 bytes	Packet Type = 81 2 bytes	hClient ID 2 bytes	Transparent Color Enable 1 byte
--------------------------	-----------------------------	-----------------------	------------------------------------

Reserved 1 1 byte	Alpha-Cursor Identifier 2 bytes	Data Format Descriptor 2 bytes	Transparent Pixel Value 4 bytes	CRC 2 bytes
----------------------	------------------------------------	-----------------------------------	------------------------------------	----------------

FIG. 33

Packet Length 2 bytes	Packet Type = 82 2 bytes	hClient ID 2 bytes	Parameter CRC 2 bytes
--------------------------	-----------------------------	-----------------------	--------------------------

Guard Time 1 64 bytes	Measurement Period 64 bytes	All Zero 2 bytes	Guard Time 2 64 bytes
--------------------------	--------------------------------	---------------------	--------------------------

FIG. 34

20/86




FIG. 35

21/86

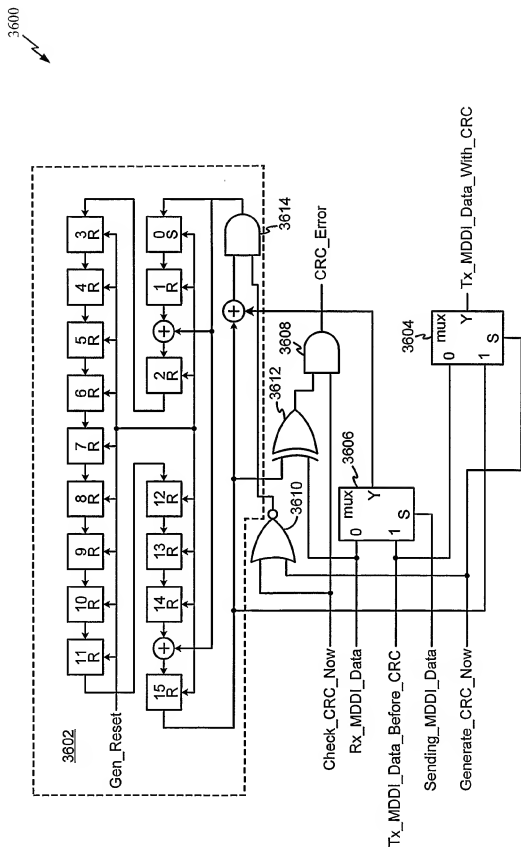


FIG. 36

22/86




FIG. 37A




FIG. 37B

23/86




FIG. 38




FIG. 39

24/86




FIG. 40

25/86


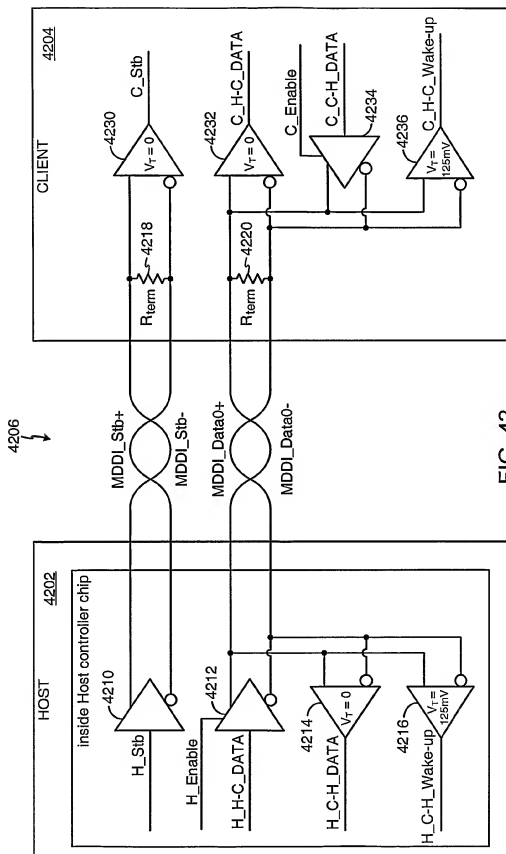



FIG. 41

26/86



27/86

FIG. 43A



Link Entering Hibernation Timing




FIG. 43B

Host-Initiated Wake-up Sequence




FIG. 43C

Client-Initiated Wake-up Sequence

28/86




FIG. 44





FIG. 45

29/86



30/86




FIG. 47

31/86

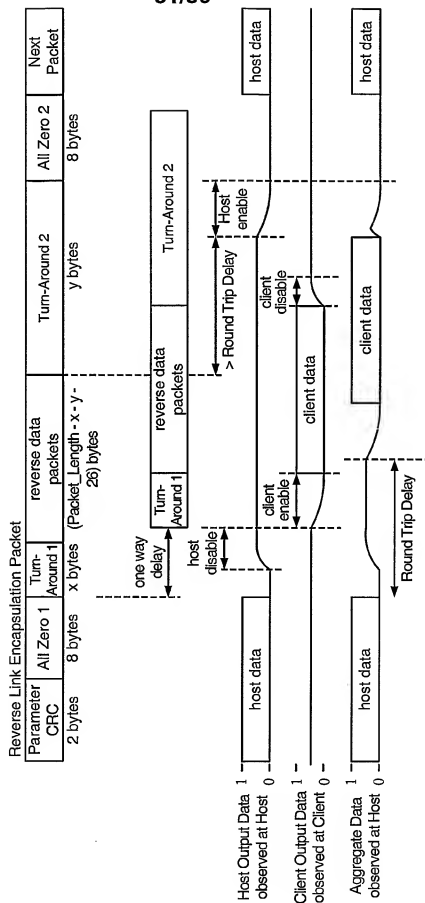


FIG. 48

32/86

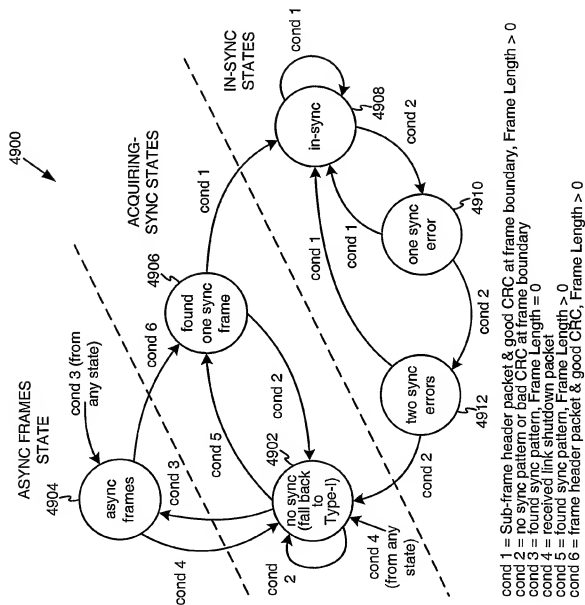


FIG. 49

33/86




FIG. 50

34/86




FIG. 51

35/86




FIG. 52A

36/86




FIG. 52B

37/86




FIG. 53

38/86

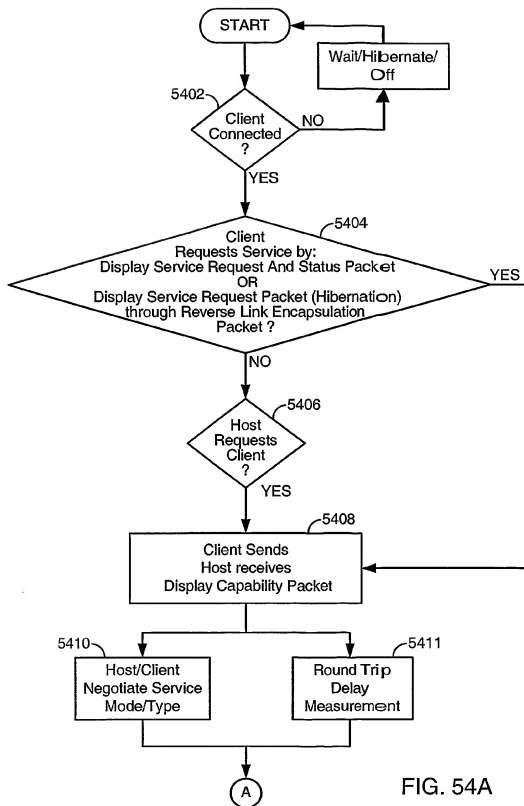


FIG. 54A

39/86

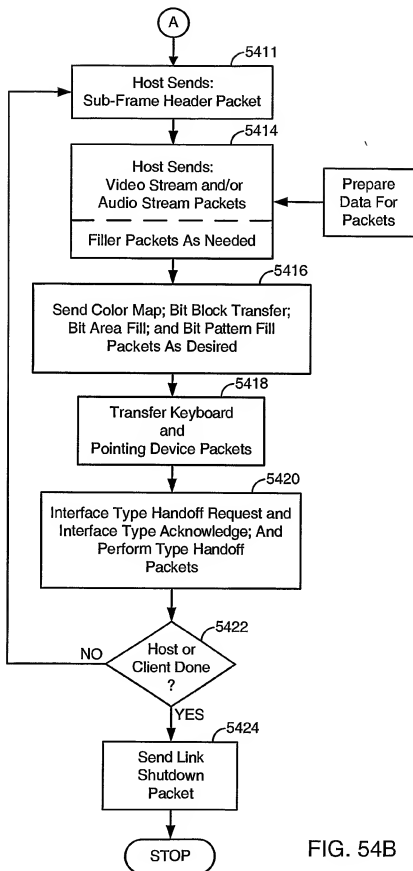


FIG. 54B

40/86




FIG. 55

41/86

Forward Link Skew Calibration Packet

Packet Length	Packet Type = 83	hClient ID	Parameter CRC	All Zero 1	Calibration Data Sequence	All Zero 2
2 bytes	2 bytes	2 bytes	2 bytes	8 bytes	Packet Length – 22 bytes	8 bytes

FIG. 56

42/86

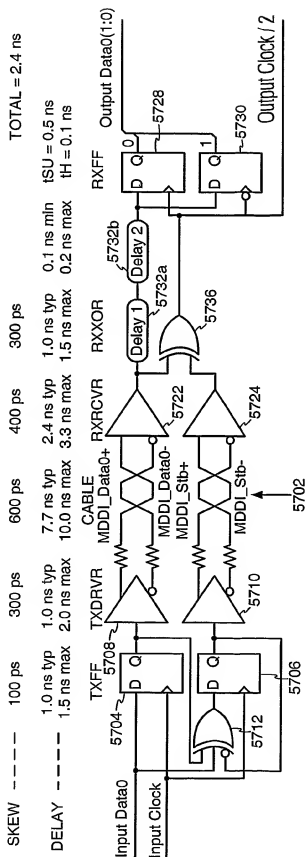


FIG. 57

43/86

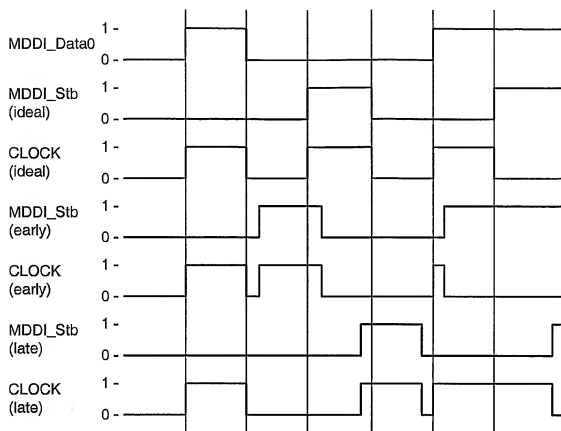


FIG. 58

44/86

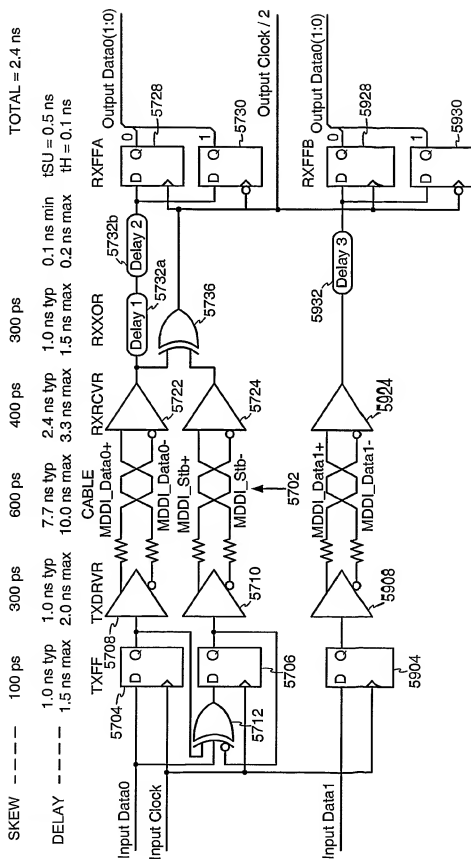


FIG. 59

45/86




FIG. 60A




FIG. 60B

47/86




FIG. 60C

48/86




FIG. 61

49/86

Type-I Data Example

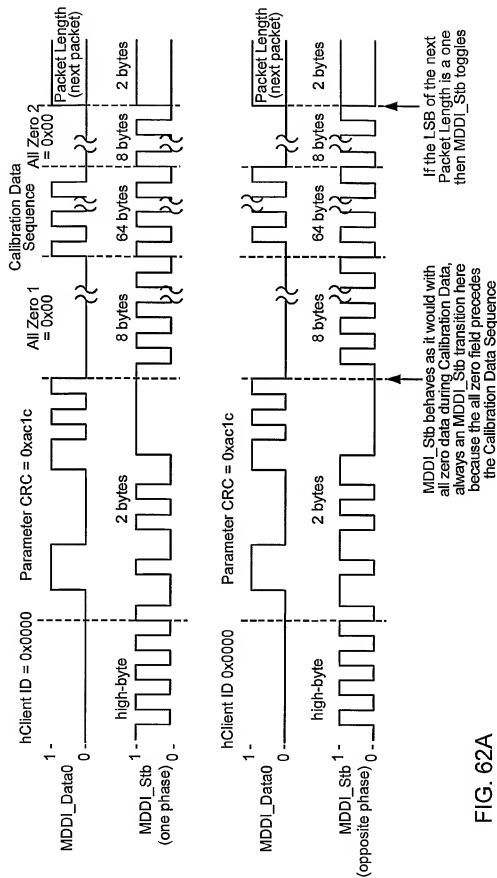


FIG. 62A

50/86

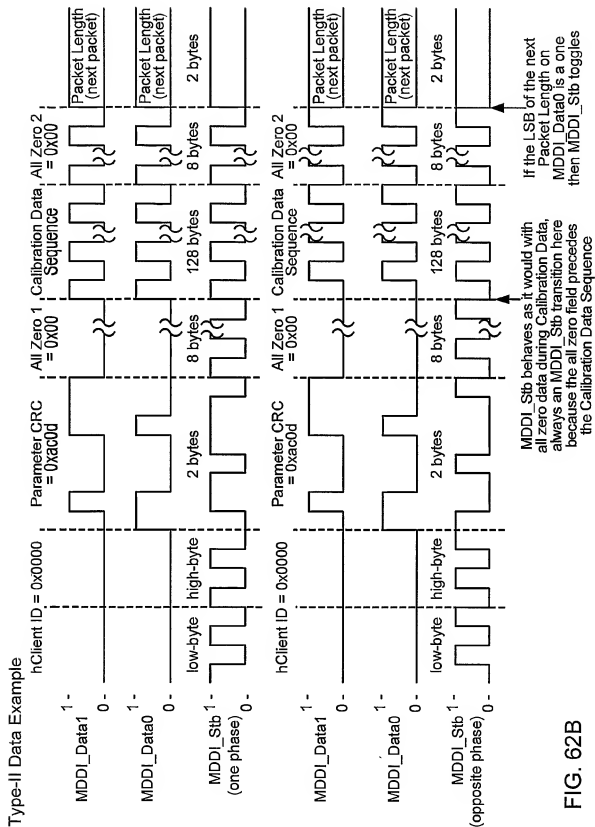


FIG. 62B

51/86

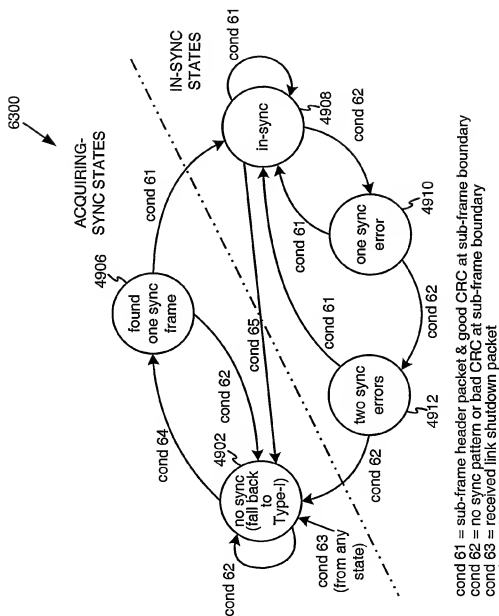


FIG. 63

52/86




FIG. 64

53/86




FIG. 65

54/86

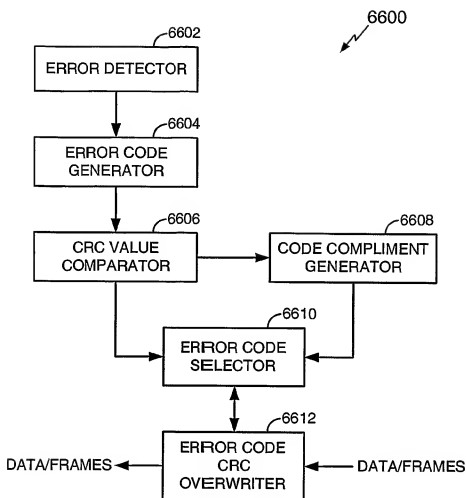


FIG. 66

55/86

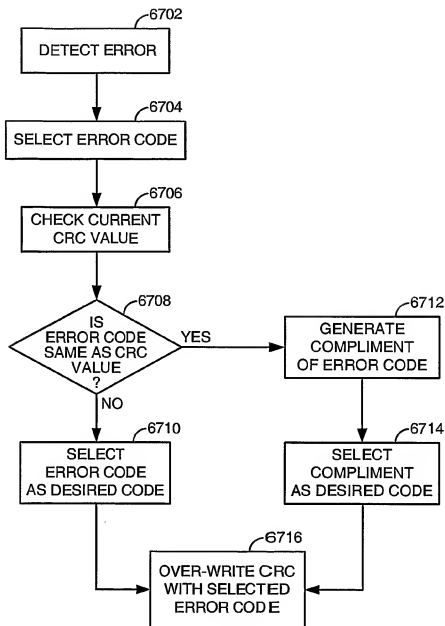


FIG. 67A

56/86

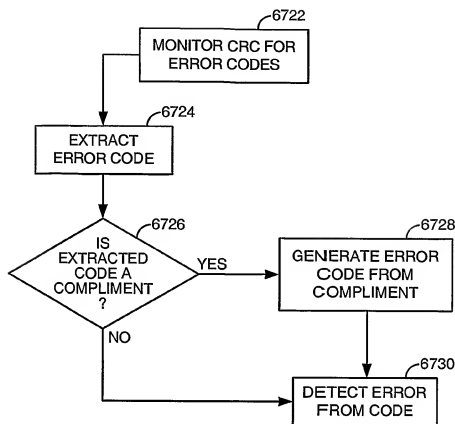






FIG. 67B

57/86



58/86




59/86

VCP Feature Reply Packet

Packet Length	Packet Type = 129	cClient ID	MCCS Version	Reply Sequence Number	Number of Features in List	VCP Feature Reply List	CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	(Packet_Length - 12 bytes)	2 bytes

FIG. 70

60/86




61/86

FIG. 72

Set VCP Feature Packet

Packet Length	Packet Type = 130	hClient ID	MCCS VCP Code	Number of Values in List	Control Value List	CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	(Packet_Length - 10 bytes)	2 bytes

62/86



63/86

Valid Parameter Reply Packet							
Packet Length	Packet Type = 132	cClient ID	MCCS VCP Code	Response Code	Reply Sequence Number	Number of Values in List	VCP Parameter Reply List
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	(Packet_Length - 14 bytes) 2 bytes

FIG. 74

64/86

Scaled Video Stream Capability Packet

Packet Length	Packet Type = 143	cClient ID	Max Number of Streams	Source Max X Size	Source Max Y Size
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

RGB Capability	Monochrome Capability	Reserved 1	Y Cr Cb Capability	Capability Bits	Reserved 2	CRC
2 bytes	1 byte	1 byte	2 bytes	1 byte	1 byte	2 bytes

FIG. 75

65/86

Scaled Video Stream Setup Packet

Packet Length	Packet Type = 136	hClient ID	Stream ID	Video Data Format Descriptor	Pixel Data Attributes
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

X Left Edge	Y Top Edge	X Right Edge	Y Bottom Edge	X Image Size	Y Image Size	CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes

FIG. 76

66/86




FIG. 77


Scaled Video Stream Packet

Packet Length	Packet Type = 18	hClient ID	Stream ID	Pixel Data Attributes	Pixel Count	Parameter CRC	Pixel Data	Pixel Data CRC
2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	2 bytes	Packet Length - 12 bytes	2 bytes

67/86

FIG. 78

68/86



69/86

Valid Status Reply List Packet				
Packet Length 2 bytes	Packet Type = 139 2 bytes	oClient ID 2 bytes	Number of Values in List 2 bytes	Valid Parameter Reply List Packet Length - 8 bytes 2 bytes
				CRC 2 bytes

FIG. 80

70/86

Personal Display Capability Packet

Packet Length	2 bytes	Packet Type = 141	cClient ID	2 bytes	Sub-Pixel Layout	1 byte	Pixel Shape	1 byte	Horizontal Field of View	1 byte	Vertical Field of View	1 byte	Visual Axis Crossing	1 byte
Lft./Rt. Image Overlap	1 byte	See Through	1 byte	Maximum Brightness	1 byte	Optical Capability	2 bytes	Minimum IPD	1 byte	Maximum IPD	1 byte	Points of Field Curvature list (25 2-Byte Values)	50 bytes	CRC
														2 bytes

71/86




FIG. 82


72/86

Client Error Report Packet

Packet Length	Packet Type = 142	cClient ID	Number of List Items	Error Code List	CRC
2 bytes	2 bytes	2 bytes	2 bytes	Packet Length - 8 bytes	2 bytes

FIG. 83

73/86



74/86

Client Identification Packet

Packet Length 2 bytes	Packet Type = 144 2 bytes	cClient ID 2 bytes	Week of Mfr 1 byte	Year of Mfr 1 byte	Length of Mfr Name 2 bytes	Length of Product Name 2 bytes	Length of Serial Number 2 bytes
Manufacturer Name String Length of Mfr Name bytes		Product Name String Length of Mfr Name bytes		Serial Number String Length of Mfr Name bytes		CRC 2 bytes	

FIG. 85

75/86

Alternate Display Capability Packet

Packet Length	Packet Type = 145	cClient ID	Alt Display Number	Reserved 1	Bitmap Width	Bitmap Height	Display Window Width	Display Window Height
2 bytes	2 bytes	2 bytes	1 byte	1 byte	2 bytes	2 bytes	2 bytes	2 bytes

Color Map RGB Width	RGB Capability	Monochrome Capability	Reserved 2	Y Cb Cr Capability	Display Feature Capability	Reserved 3	CRC
2 bytes	2 bytes	1 byte	1 byte	2 bytes	1 byte	1 byte	2 bytes

FIG. 86

76/86

Register Access Packet							
Packet Length	Packet Type = 146	bClient ID	Read/Write Flags	Register Address	Parameter CRC	Register Data List	CRC
2 bytes	2 bytes	2 bytes	2 bytes	4 bytes	2 bytes	Packet Length - 14 bytes	2 bytes

FIG. 87

77/86




FIG. 88A




FIG. 88B




FIG. 88C

78/86

Two Buffers, display refresh faster than Image transfer

Image Transfer		Image Transfer		Image Transfer		Image Transfer	
frame 1, buffer A 20% 40% 60% 80%	frame 2, buffer B 20% 40% 60% 80%	frame 3, buffer A 20% 40% 60% 80%	frame 4, buffer B 20% 40% 60% 80%	frame 5, buffer A 20% 40% 60% 80%	frame 6, buffer B 20% 40% 60% 80%		
Display Refresh		Display Refresh		Display Refresh		Display Refresh	
frame 0, buf B 20% 40% 60% 80%	frame 0, buf B 20% 40% 60% 80%	frame 1, buf A 20% 40% 60% 80%	frame 2, buf B 20% 40% 60% 80%	frame 3, buf A 20% 40% 60% 80%	frame 4, buf B 20% 40% 60% 80%	frame 5, buf A 20% 40% 60% 80%	frame 6, buf B 20% 40% 60% 80%

Two Buffers, display refresh much faster than image transfer

Image Transfer		Image Transfer		Image Transfer		Image Transfer	
frame 1, buffer A 20% 40% 60% 80%	frame 2, buffer B 20% 40% 60% 80%	frame 3, buffer A 20% 40% 60% 80%	frame 4, buffer B 20% 40% 60% 80%	frame 5, buffer A 20% 40% 60% 80%	frame 6, buffer B 20% 40% 60% 80%		
Display Refresh		Display Refresh		Display Refresh		Display Refresh	
fm 0, b=B 20 40 60 80 % % % %	fm 1, b=A 20 40 60 80 % % % %	fm 2, b=B 20 40 60 80 % % % %	fm 3, b=A 20 40 60 80 % % % %	fm 4, b=B 20 40 60 80 % % % %	fm 5, b=A 20 40 60 80 % % % %	fm 6, b=B 20 40 60 80 % % % %	fm 7, b=A 20 40 60 80 % % % %

FIG. 89

79/86

Two Buffers, display refresh slower than Image transfer




FIG. 90

80/86

Image Transfer . Two Buffers, display refresh much faster than Image transfer, small video window

frame 1, buffer A 20% 40% 60% 80%	frame 2, buffer B 20% 40% 60% 80%	frame 3, buffer A 20% 40% 60% 80%	frame 4, buffer B 20% 40% 60% 80%	frame 5, buffer A 20% 40% 60% 80%	frame 6, buffer B 20% 40% 60% 80%
fm 0, b=B fm 0, b=B fm 0, b=B fm 0, b=B fm 0, b=B fm 0, b=B	fm 1, b=A fm 1, b=A fm 1, b=A fm 1, b=A fm 1, b=A fm 1, b=A	fm 2, b=B fm 2, b=B fm 2, b=B fm 2, b=B fm 2, b=B fm 2, b=B	fm 3, b=A fm 3, b=A fm 3, b=A fm 3, b=A fm 3, b=A fm 3, b=A	fm 4, b=B fm 4, b=B fm 4, b=B fm 4, b=B fm 4, b=B fm 4, b=B	fm 5, b=A fm 5, b=A fm 5, b=A fm 5, b=A fm 5, b=A fm 5, b=A
20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80	20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80	20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80	20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80	20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80	20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80 20 40 60 80
% %	% %	% %	% %	% %	% %

FIG. 91

81/86

Three Buffers, display refresh much faster than Image transfer, any-size video window

Image Transfer

frame 1, buffer A 20% 40% 60% 80%	frame 2, buffer B 20% 40% 60% 80%	frame 3, buffer C 20% 40% 60% 80%	frame 4, buffer A 20% 40% 60% 80%	frame 5, buffer B 20% 40% 60% 80%	frame 6, buffer C 20% 40% 60% 80%
--------------------------------------	--------------------------------------	--------------------------------------	--------------------------------------	--------------------------------------	--------------------------------------

Display Refresh

fm 0, b=C	fm 0, b=C	fm 1, b=A	fm 1, b=A	fm 2, b=B	fm 2, b=B	fm 3, b=C	fm 3, b=C	fm 4, b=A	fm 4, b=A	fm 5, b=B	fm 5, b=B	fm 6, b=C	fm 6, b=C
20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80	20 40 60 80
% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %	% % % %

FIG. 92

82/86

Three Buffers, display refresh slower than Image transfer



Image Transfer				
frame 1, buffer A 20% 40% 60% 80%	frame 2, buffer B 20% 40% 60% 80%	frame 3, buffer C 20% 40% 60% 80%	frame 4, buffer A 20% 40% 60% 80%	frame 5, buffer B 20% 40% 60% 80%
				frame 6, buffer C 20% 40% 60% 80%
				frame 7 20%
Display Refresh				
frame 0, buffer C 20% 40% 60% 80%	frame 1 & 4, buffer A 20% 40% 60% 80%	 60%	frame 3, buffer C 20% 40% 60% 80%	frame 4 & 7, buffer A 20% 40% 60% 80%
				 60%

FIG. 93

83/86

One Buffer, display refresh faster than Image transfer

Image Transfer		frame 1 20% 40% 60% 80%	frame 2 20% 40% 60% 80%	frame 3 20% 40% 60% 80%	frame 4 20% 40% 60% 80%	frame 5 20% 40% 60% 80%	frame 6 20% 40% 60% 80%
Display Refresh		frame 0 20% 40% 60% 80%	frame 1 20% 40% 60% 80%	frame 2 20% 40% 60% 80%	frame 3 20% 40% 60% 80%	frame 4 20% 40% 60% 80%	frame 5 20% 40% 60% 80%
		frame 0 20% 40% 60% 80%	frame 2 20% 40% 60% 80%	frame 3 & 2 20% 40% 60% 80%	frame 3 20% 40% 60% 80%	frame 4 20% 40% 60% 80%	frame 5 20% 40% 60% 80%
		frame 0 20% 40% 60% 80%	frame 1 20% 40% 60% 80%	frame 2 20% 40% 60% 80%	frame 3 20% 40% 60% 80%	frame 4 20% 40% 60% 80%	frame 5 20% 40% 60% 80%

One Buffer, display refresh much faster than image transfer

Image Transfer		frame 1 20% 40% 60% 80%	frame 2 20% 40% 60% 80%	frame 3 20% 40% 60% 80%	frame 4 20% 40% 60% 80%	frame 5 20% 40% 60% 80%	frame 6 20% 40% 60% 80%
Display Refresh		fm 0 20% 40% 60% 80%	fm 1 20% 40% 60% 80%	fm 2 20% 40% 60% 80%	fm 3 20% 40% 60% 80%	fm 4 20% 40% 60% 80%	fm 5 20% 40% 60% 80%
		fm 0 20% 40% 60% 80%	fm 1 20% 40% 60% 80%	fm 2 20% 40% 60% 80%	fm 3 20% 40% 60% 80%	fm 4 20% 40% 60% 80%	fm 5 20% 40% 60% 80%
		fm 0 20% 40% 60% 80%	fm 1 20% 40% 60% 80%	fm 2 20% 40% 60% 80%	fm 3 20% 40% 60% 80%	fm 4 20% 40% 60% 80%	fm 5 20% 40% 60% 80%

FIG. 94

84/86





FIG. 95

85/86




FIG. 96




FIG. 97

INTERNATIONAL SEARCH REPORT

 Int'l Application No
 PCT/US2005/008832

 A. CLASSIFICATION OF SUBJECT MATTER
 IPC 7 H04L29/06 H04L12/28 H04L12/56

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 H04L

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, IBM-TDB, WPI Data, INSPEC, COMPENDEX

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category *	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No
X	WO 03/023587 A2 (QUALCOMM, INCORPORATED) 20 March 2003 (2003-03-20) the whole document	1-33
A	VESA: "VESA Plug and Display Standard Version 1.0" VESA, 11 June 1997 (1997-06-11), XP002123075 VESA the whole document	1-33
A	WO 96/42158 A (TELEFONAKTIEBOLAGET LM ERICSSON; PETERSEN, LARS, GOERAN; KUNDEL, MIKAE) 27 December 1996 (1996-12-27) abstract page 1, line 1 - page 5, line 1 page 6, line 8 - page 25, line 25 figures 1-14	1-33

☐ Further documents are listed in the continuation of box C

☒ Patent family members are listed in annex

* Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

S document member of the same patent family

Date of the actual completion of the international search

30 August 2005

Date of mailing of the international search report

07/09/2005

Name and mailing address of the ISA

 European Patent Office, P.B. 5818 Patentlaan 2
 NL - 2280 HV Rijswijk
 Tel (+31-70) 340-2040, Tx 31 651 epo nl,
 Fax (+31-70) 340-3016

Authorized officer

Körbler, G

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No

PCT/US2005/008832

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
WO 03023587	A2	20-03-2003	US 2003033417 A1 13-02-2003
			CA 2459941 A1 20-03-2003
			CN 1575448 A 02-02-2005
			EP 1423778 A2 02-06-2004
			US 2004199652 A1 07-10-2004
WO 9642158	A	27-12-1996	SE 506540 C2 12-01-1998
			AU 6021896 A 09-01-1997
			BR 9608762 A 06-07-1999
			CA 2224196 A1 27-12-1996
			CN 1192836 A ,C 09-09-1998
			DE 69634048 D1 20-01-2005
			EP 0872085 A1 21-10-1998
			JP 11507788 T 06-07-1999
			RU 2156035 C2 10-09-2000
			SE 9502142 A 14-12-1996
			WO 9642158 A1 27-12-1996
			US 5963564 A 05-10-1999